

MATRIX CONCENTRATION INEQUALITIES VIA THE METHOD OF EXCHANGEABLE PAIRS

LESTER MACKEY AND MICHAEL I. JORDAN*
RICHARD Y. CHEN, BRENDAN FARRELL, AND JOEL A. TROPP†

ABSTRACT. This paper derives exponential concentration inequalities and polynomial moment inequalities for the spectral norm of a random matrix. The analysis requires a matrix extension of the scalar concentration theory developed by Sourav Chatterjee using Stein’s method of exchangeable pairs. When applied to a sum of independent random matrices, this approach yields matrix generalizations of the classical inequalities due to Hoeffding, Bernstein, Khintchine, and Rosenthal. The same technique delivers bounds for sums of dependent random matrices and more general matrix-valued functions of dependent random variables.

This paper is based on two independent manuscripts from mid-2011 that both applied the method of exchangeable pairs to establish matrix concentration inequalities. One manuscript is by Mackey and Jordan; the other is by Chen, Farrell, and Tropp. The authors have combined this research into a single unified presentation, with equal contributions from both groups.

1. INTRODUCTION

Matrix concentration inequalities control the fluctuations of a random matrix about its mean. At present, these results provide an effective method for studying sums of independent random matrices and matrix martingales [Oli09, Tro11a, Tro11b, HKZ11]. They have been used to streamline the analysis of structured random matrices in a range of applications, including statistical estimation [Kol11], randomized linear algebra [Git11, CD11b], stability of least-squares approximation [CDL11], combinatorial and robust optimization [So11, CSW11], matrix completion [Gro11, Rec11, MTJ11], and random graph theory [Oli09]. These works comprise only a small sample of the papers that rely on matrix concentration inequalities. Nevertheless, it remains common to encounter new classes of random matrices that we cannot treat with the available techniques.

The purpose of this paper is to lay the foundations of a new approach for analyzing structured random matrices. Our work is based on Chatterjee’s technique for developing scalar concentration inequalities [Cha07, Cha08] via Stein’s method of exchangeable pairs [Ste72]. We extend this argument to the matrix setting, where we use it to establish exponential concentration results (Theorems 4.1 and 5.1) and polynomial moment inequalities (Theorem 7.1) for the spectral norm of a random matrix.

To illustrate the power of this idea, we show that our general results imply several important concentration bounds for a sum of independent, random, Hermitian matrices [LPP91, JX03, Tro11b]. We obtain a matrix Hoeffding inequality with optimal constants (Corollary 4.2) and a version of the matrix Bernstein inequality (Corollary 5.2). Our techniques also yield concise proofs of the matrix Khintchine inequality (Corollary 7.4) and the matrix Rosenthal inequality (Corollary 7.5).

The method of exchangeable pairs also applies to matrices constructed from dependent random variables. We offer a hint of the prospects by establishing concentration results for several other

Date: 27 January 2012.

*Department of Electrical Engineering and Computer Sciences and Department of Statistics, University of California, Berkeley.

†Department of Computing and Mathematical Sciences, California Institute of Technology.

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 27 JAN 2012		2. REPORT TYPE		3. DATES COVERED 00-00-2012 to 00-00-2012	
4. TITLE AND SUBTITLE Matrix Concentration Inequalities via the Method of Exchangeable Pairs				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) California Institute of Technology, Department of Computing and Mathematical Sciences, Pasadena, CA, 91125				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT This paper derives exponential concentration inequalities and polynomial moment inequalities for the spectral norm of a random matrix. The analysis requires a matrix extension of the scalar concentration theory developed by Sourav Chatterjee using Stein's method of exchangeable pairs. When applied to a sum of independent random matrices, this approach yields matrix generalizations of the classical inequalities due to Hoeffding, Bernstein, Khintchine, and Rosenthal. The same technique delivers bounds for sums of dependent random matrices and more general matrix-valued functions of dependent random variables.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 29	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

classes of random matrices. In Section 9, we consider sums of dependent matrices that satisfy a conditional zero-mean property. In Section 10, we treat a broad class of combinatorial matrix statistics. Finally, in Section 11, we analyze general matrix-valued functions that have a self-reproducing property.

1.1. Notation and Preliminaries. The symbol $\|\cdot\|$ is reserved for the spectral norm, which returns the largest singular value of a general complex matrix.

We write \mathbb{M}^d for the algebra of all $d \times d$ complex matrices. The trace and normalized trace of a square matrix are defined as

$$\mathrm{tr} \mathbf{B} := \sum_{j=1}^d b_{jj} \quad \text{and} \quad \bar{\mathrm{tr}} \mathbf{B} := \frac{1}{d} \sum_{j=1}^d b_{jj} \quad \text{for } \mathbf{B} \in \mathbb{M}^d.$$

We define the linear space \mathbb{H}^d of Hermitian $d \times d$ matrices. *All matrices in this paper are Hermitian unless explicitly stated otherwise.* The symbols $\lambda_{\max}(\mathbf{A})$ and $\lambda_{\min}(\mathbf{A})$ refer to the algebraic maximum and minimum eigenvalues of a matrix $\mathbf{A} \in \mathbb{H}^d$. For each interval $I \subset \mathbb{R}$, we define the set of Hermitian matrices whose eigenvalues fall in that interval:

$$\mathbb{H}^d(I) := \{\mathbf{A} \in \mathbb{H}^d : \lambda_{\max}(\mathbf{A}), \lambda_{\min}(\mathbf{A}) \in I\}.$$

The set \mathbb{H}_+^d consists of all positive-semidefinite (psd) $d \times d$ matrices. Curly inequalities refer to the semidefinite partial order on Hermitian matrices. For example, we write $\mathbf{A} \preceq \mathbf{B}$ to signify that the matrix $\mathbf{B} - \mathbf{A}$ is psd.

We require operator convexity properties of the matrix square so often that we state them now.

$$\left(\frac{\mathbf{A} + \mathbf{B}}{2}\right)^2 \preceq \frac{\mathbf{A}^2 + \mathbf{B}^2}{2} \quad \text{for all } \mathbf{A}, \mathbf{B} \in \mathbb{H}^d. \quad (1.1)$$

More generally, we have the operator Jensen inequality

$$(\mathbb{E} \mathbf{X})^2 \preceq \mathbb{E} \mathbf{X}^2, \quad (1.2)$$

valid for any random Hermitian matrix, provided that $\mathbb{E} \|\mathbf{X}\|^2 < \infty$. To verify this result, simply expand the inequality $\mathbb{E}(\mathbf{X} - \mathbb{E} \mathbf{X})^2 \succeq \mathbf{0}$. The operator Jensen inequality also holds for conditional expectation, again provided that $\mathbb{E} \|\mathbf{X}\|^2 < \infty$.

2. EXCHANGEABLE PAIRS OF RANDOM MATRICES

Our approach to studying random matrices is based on the method of exchangeable pairs, which originates in the work of Charles Stein [Ste72] on normal approximation for a sum of dependent random variables. In this section, we explain how some central ideas from this theory extend to matrices.

2.1. Matrix Stein Pairs. We begin with the definition of an exchangeable pair.

Definition 2.1 (Exchangeable Pair). Let Z and Z' be random variables taking values in a Polish space¹ \mathcal{Z} . We say that (Z, Z') is an *exchangeable pair* if it has the same distribution as (Z', Z) . In particular, Z and Z' must share the same distribution.

We can obtain a lot of information about the fluctuations of a random matrix \mathbf{X} if we can construct a good exchangeable pair $(\mathbf{X}, \mathbf{X}')$. With this motivation in mind, let us introduce a special class of exchangeable pairs.

¹A topological space is *Polish* if we can equip it with a metric to form a complete, separable metric space.

Definition 2.2 (Matrix Stein Pair). Let (Z, Z') be an exchangeable pair of random variables taking values in a Polish space \mathcal{Z} , and let $\Psi : \mathcal{Z} \rightarrow \mathbb{H}^d$ be a measurable function. Define the random Hermitian matrices

$$\mathbf{X} := \Psi(Z) \quad \text{and} \quad \mathbf{X}' := \Psi(Z').$$

We say that $(\mathbf{X}, \mathbf{X}')$ is a *matrix Stein pair* if there is a constant $\alpha \in (0, 1]$ for which

$$\mathbb{E}[\mathbf{X} - \mathbf{X}' | Z] = \alpha \mathbf{X} \quad \text{almost surely.} \quad (2.1)$$

The constant α is called the *scale factor* of the pair. When discussing a matrix Stein pair $(\mathbf{X}, \mathbf{X}')$, we always assume that $\mathbb{E} \|\mathbf{X}\|^2 < \infty$.

A matrix Stein pair $(\mathbf{X}, \mathbf{X}')$ has several useful properties. First, $(\mathbf{X}, \mathbf{X}')$ always forms an exchangeable pair. Second, it must be the case that $\mathbb{E} \mathbf{X} = \mathbf{0}$. Indeed,

$$\mathbb{E} \mathbf{X} = \frac{1}{\alpha} \mathbb{E} [\mathbb{E}[\mathbf{X} - \mathbf{X}' | Z]] = \frac{1}{\alpha} \mathbb{E}[\mathbf{X} - \mathbf{X}'] = \mathbf{0}$$

because of the identity (2.1), the tower property of conditional expectation, and the exchangeability of $(\mathbf{X}, \mathbf{X}')$. In Section 2.4, we construct a matrix Stein pair for a sum of centered, independent random matrices. More sophisticated examples appear in Sections 9, 10, and 11.

2.2. The Method of Exchangeable Pairs. A well-chosen matrix Stein pair $(\mathbf{X}, \mathbf{X}')$ provides a surprisingly powerful tool for studying the random matrix \mathbf{X} . The technique depends on a simple but fundamental technical lemma.

Lemma 2.3 (Method of Exchangeable Pairs). *Suppose that $(\mathbf{X}, \mathbf{X}') \in \mathbb{H}^d \times \mathbb{H}^d$ is a matrix Stein pair with scale factor α . Let $\mathbf{F} : \mathbb{H}^d \rightarrow \mathbb{H}^d$ be a measurable function that satisfies the regularity condition*

$$\mathbb{E} \|(\mathbf{X} - \mathbf{X}') \cdot \mathbf{F}(\mathbf{X})\| < \infty. \quad (2.2)$$

Then

$$\mathbb{E} [\mathbf{X} \cdot \mathbf{F}(\mathbf{X})] = \frac{1}{2\alpha} \mathbb{E} [(\mathbf{X} - \mathbf{X}')(\mathbf{F}(\mathbf{X}) - \mathbf{F}(\mathbf{X}'))]. \quad (2.3)$$

To appreciate this result, recall that we can characterize the distribution of a random matrix by integrating it against a sufficiently large class of test functions. The additional randomness in the Stein pair furnishes an alternative expression for the expected product of \mathbf{X} and the test function \mathbf{F} . The identity (2.3) is valuable because it allows us to estimate this integral using the smoothness properties of the function \mathbf{F} and the discrepancy between \mathbf{X} and \mathbf{X}' .

Proof. Suppose $(\mathbf{X}, \mathbf{X}')$ is a matrix Stein pair constructed from an auxiliary exchangeable pair (Z, Z') . The defining property (2.1) of the Stein pair implies

$$\alpha \cdot \mathbb{E}[\mathbf{X} \cdot \mathbf{F}(\mathbf{X})] = \mathbb{E} [\mathbb{E}[\mathbf{X} - \mathbf{X}' | Z] \cdot \mathbf{F}(\mathbf{X})] = \mathbb{E}[(\mathbf{X} - \mathbf{X}') \mathbf{F}(\mathbf{X})].$$

We have used the regularity condition (2.2) to invoke the pull-through property of conditional expectation. Since $(\mathbf{X}, \mathbf{X}')$ is an exchangeable pair,

$$\alpha \cdot \mathbb{E}[\mathbf{X} \cdot \mathbf{F}(\mathbf{X})] = \mathbb{E}[(\mathbf{X} - \mathbf{X}') \mathbf{F}(\mathbf{X})] = \mathbb{E}[(\mathbf{X}' - \mathbf{X}) \mathbf{F}(\mathbf{X}')] = -\mathbb{E}[(\mathbf{X} - \mathbf{X}') \mathbf{F}(\mathbf{X}')].$$

The identity (2.3) follows when we average the two preceding displays. \square

2.3. The Conditional Variance. To each matrix Stein pair $(\mathbf{X}, \mathbf{X}')$, we may associate a random matrix called the *conditional variance* of \mathbf{X} . The ultimate purpose of this paper is to argue that the spectral norm of \mathbf{X} is unlikely to be large when the conditional variance is small.

Definition 2.4 (Conditional Variance). Suppose that $(\mathbf{X}, \mathbf{X}')$ is a matrix Stein pair, constructed from an auxiliary exchangeable pair (Z, Z') . The *conditional variance* is the random matrix

$$\Delta_{\mathbf{X}} := \Delta_{\mathbf{X}}(Z) := \frac{1}{2\alpha} \mathbb{E}[(\mathbf{X} - \mathbf{X}')^2 | Z], \quad (2.4)$$

where α is the scale factor of the pair. We may take any version of the conditional expectation in this definition.

The conditional variance $\Delta_{\mathbf{X}}$ should be regarded as a stochastic estimate for the variance of the random matrix \mathbf{X} . Indeed,

$$\mathbb{E}[\Delta_{\mathbf{X}}] = \mathbb{E} \mathbf{X}^2. \quad (2.5)$$

This identity follows immediately from Lemma 2.3 with the choice $\mathbf{F}(\mathbf{X}) = \mathbf{X}$.

2.4. Example: A Sum of Independent Random Matrices. To make the definitions in this section more vivid, we describe a simple but important example of a matrix Stein pair. Consider an independent sequence $Z := (\mathbf{Y}_1, \dots, \mathbf{Y}_n)$ of random Hermitian matrices that satisfy $\mathbb{E} \mathbf{Y}_k = \mathbf{0}$ and $\mathbb{E} \|\mathbf{Y}_k\|^2 < \infty$ for each k . Introduce the random series

$$\mathbf{X} := \mathbf{Y}_1 + \dots + \mathbf{Y}_n.$$

Let us explain how to build a good matrix Stein pair $(\mathbf{X}, \mathbf{X}')$. We need the exchangeable counterpart \mathbf{X}' to have the same distribution as \mathbf{X} , but it should also be close to \mathbf{X} so that we can control the conditional variance. To achieve these goals, we construct \mathbf{X}' by picking a summand from \mathbf{X} at random and replacing it with a fresh copy.

Formally, let \mathbf{Y}'_k be an independent copy of \mathbf{Y}_k for each index k , and draw a random index K uniformly at random from $\{1, \dots, n\}$ and independently from everything else. Define the random sequence $Z' := (\mathbf{Y}_1, \dots, \mathbf{Y}_{K-1}, \mathbf{Y}'_K, \mathbf{Y}_{K+1}, \dots, \mathbf{Y}_n)$. It is easy to check that (Z, Z') forms an exchangeable pair. The random matrix

$$\mathbf{X}' := \mathbf{Y}_1 + \dots + \mathbf{Y}_{K-1} + \mathbf{Y}'_K + \mathbf{Y}_{K+1} + \dots + \mathbf{Y}_n$$

is thus an exchangeable counterpart for \mathbf{X} . To verify that $(\mathbf{X}, \mathbf{X}')$ is a Stein pair, calculate that

$$\mathbb{E}[\mathbf{X} - \mathbf{X}' | Z] = \mathbb{E}[\mathbf{Y}_K - \mathbf{Y}'_K | Z] = \frac{1}{n} \sum_{k=1}^n \mathbb{E}[\mathbf{Y}_k - \mathbf{Y}'_k | Z] = \frac{1}{n} \sum_{k=1}^n \mathbf{Y}_k = \frac{1}{n} \mathbf{X}.$$

The third identity holds because \mathbf{Y}'_k is a centered random matrix that is independent from Z . Therefore, $(\mathbf{X}, \mathbf{X}')$ is a matrix Stein pair with scale factor $\alpha = n^{-1}$.

Next, we compute the conditional variance.

$$\begin{aligned} \Delta_{\mathbf{X}} &= \frac{n}{2} \cdot \mathbb{E}[(\mathbf{X} - \mathbf{X}')^2 | Z] \\ &= \frac{n}{2} \cdot \frac{1}{n} \sum_{k=1}^n \mathbb{E}[(\mathbf{Y}_k - \mathbf{Y}'_k)^2 | Z] \\ &= \frac{1}{2} \sum_{k=1}^n [\mathbf{Y}_k^2 - \mathbf{Y}_k(\mathbb{E} \mathbf{Y}'_k) - (\mathbb{E} \mathbf{Y}'_k) \mathbf{Y}_k + \mathbb{E}(\mathbf{Y}'_k)^2] \\ &= \frac{1}{2} \sum_{k=1}^n (\mathbf{Y}_k^2 + \mathbb{E} \mathbf{Y}_k^2). \end{aligned} \quad (2.6)$$

For the third relation, expand the square and invoke the pull-through property of conditional expectation. We may drop the conditioning because \mathbf{Y}'_k is independent from Z . In the last line, we apply the property that \mathbf{Y}'_k has the same distribution as \mathbf{Y}_k .

The expression (2.6) shows that we can control the size of the conditional expectation uniformly if we can control the size of the individual summands. This example also teaches us that we may use the symmetries of the distribution of the random matrix to construct a matrix Stein pair.

3. EXPONENTIAL MOMENTS AND EIGENVALUES OF A RANDOM MATRIX

Our main goal in this paper is to study the behavior of the extreme eigenvalues of a random Hermitian matrix. In Section 3.2, we describe an approach to this problem that parallels the classical Laplace transform method for scalar random variables. The adaptation to the matrix setting leads us to consider the *trace* of the moment generating function (mgf) of a random matrix. After presenting this background, we explain how the method of exchangeable pairs can be used to control the growth of the trace mgf. This result, which appears in Section 3.5, is the key to our exponential concentration bounds for random matrices.

3.1. Standard Matrix Functions. Before entering the discussion, recall that a *standard matrix function* is obtained by applying a real function to the eigenvalues of a Hermitian matrix. The book [Hig08] provides an excellent treatment of this concept.

Definition 3.1 (Standard Matrix Function). Let $f : I \rightarrow \mathbb{R}$ be a function on an interval I of the real line. Suppose that $\mathbf{A} \in \mathbb{H}^d(I)$ has the eigenvalue decomposition $\mathbf{A} = \mathbf{Q} \cdot \text{diag}(\lambda_1, \dots, \lambda_d) \cdot \mathbf{Q}^*$ where \mathbf{Q} is a unitary matrix. Then

$$f(\mathbf{A}) := \mathbf{Q} \begin{bmatrix} f(\lambda_1) & & \\ & \ddots & \\ & & f(\lambda_d) \end{bmatrix} \mathbf{Q}^*.$$

The *spectral mapping theorem* states that $f(\lambda)$ is an eigenvalue of $f(\mathbf{A})$ if and only if λ is an eigenvalue of \mathbf{A} . This fact follows immediately from Definition 3.1.

When we apply a familiar scalar function to a Hermitian matrix, we are always referring to a standard matrix function. For instance, $|\mathbf{A}|$ is the matrix absolute value, $\exp(\mathbf{A})$ is the matrix exponential, and $\log(\mathbf{A})$ is the matrix logarithm. The latter is defined only for positive matrices.

3.2. The Matrix Laplace Transform Method. Let us introduce a matrix variant of the classical moment generating function. We learned this definition from Ahlswede–Winter [AW02, App.].

Definition 3.2 (Trace Mgf). Let \mathbf{X} be a random Hermitian matrix. The (*normalized*) *trace moment generating function* of \mathbf{X} is defined as

$$m(\theta) := m_{\mathbf{X}}(\theta) := \mathbb{E} \bar{\text{tr}} e^{\theta \mathbf{X}} \quad \text{for } \theta \in \mathbb{R}.$$

We admit the possibility that the expectation may not exist for all values of θ .

Ahlswede and Winter [AW02, App.] had the insight that the classical Laplace transform method could be extended to the matrix setting by replacing the classical mgf with the trace mgf. This adaptation allows us to obtain concentration inequalities for the extreme eigenvalues of a random Hermitian matrix using methods from matrix analysis. The following proposition distills results from the papers [AW02, Oli10, Tro11b, CGT12].

Proposition 3.3 (Matrix Laplace Transform Method). Let $\mathbf{X} \in \mathbb{H}^d$ be a random matrix with normalized trace mgf $m(\theta) := \mathbb{E} \bar{\text{tr}} e^{\theta \mathbf{X}}$. For each $t \in \mathbb{R}$,

$$\mathbb{P} \{ \lambda_{\max}(\mathbf{X}) \geq t \} \leq d \cdot \inf_{\theta > 0} \exp \{ -\theta t + \log m(\theta) \}. \quad (3.1)$$

$$\mathbb{P} \{ \lambda_{\min}(\mathbf{X}) \leq t \} \leq d \cdot \inf_{\theta < 0} \exp \{ -\theta t + \log m(\theta) \}. \quad (3.2)$$

Furthermore,

$$\mathbb{E} \lambda_{\max}(\mathbf{X}) \leq \inf_{\theta > 0} \frac{1}{\theta} [\log d + \log m(\theta)]. \quad (3.3)$$

$$\mathbb{E} \lambda_{\min}(\mathbf{X}) \geq \sup_{\theta < 0} \frac{1}{\theta} [\log d + \log m(\theta)]. \quad (3.4)$$

The estimates (3.3) and (3.4) for the expectations are usually sharp up to the logarithm of the dimension. In many situations, the tail bounds (3.1) and (3.2) are reasonable for moderate t , but they tend to overestimate the probability of a large deviation. Note that, in general, we cannot dispense with the dimensional factor d . See [Tro11b, Sec. 4] for a detailed discussion of these issues. Additional inequalities for the interior eigenvalues can be established using the minimax Laplace transform method [GT11].

Proof. To establish (3.1), fix $\theta > 0$. Using Markov's inequality, we find that

$$\begin{aligned} \mathbb{P} \{ \lambda_{\max}(\mathbf{X}) \geq t \} &= \mathbb{P} \left\{ e^{\lambda_{\max}(\theta \mathbf{X})} \geq e^{\theta t} \right\} \leq e^{-\theta t} \cdot \mathbb{E} e^{\lambda_{\max}(\theta \mathbf{X})} \\ &= e^{-\theta t} \cdot \mathbb{E} \lambda_{\max}(e^{\theta \mathbf{X}}) \leq e^{-\theta t} \cdot \mathbb{E} \operatorname{tr} e^{\theta \mathbf{X}}. \end{aligned}$$

The third relation follows from the spectral mapping theorem. The final inequality holds because the trace of a positive matrix dominates its maximum eigenvalue. Identify the normalized trace mgf, and take the infimum over θ to complete the argument.

The proof of (3.2) parallels the proof of (3.1). For $\theta < 0$,

$$\mathbb{P} \{ \lambda_{\min}(\mathbf{X}) \leq t \} = \mathbb{P} \{ \theta \lambda_{\min}(\mathbf{X}) \geq \theta t \} = \mathbb{P} \{ \lambda_{\max}(\theta \mathbf{X}) \geq \theta t \}.$$

We have used the property that $-\lambda_{\min}(\mathbf{A}) = \lambda_{\max}(-\mathbf{A})$ for each Hermitian matrix \mathbf{A} . The remainder of the argument is the same as in the preceding paragraph.

For the expectation bound (3.3), fix $\theta > 0$. Jensen's inequality yields

$$\mathbb{E} \lambda_{\max}(\mathbf{X}) = \theta^{-1} \mathbb{E} \lambda_{\max}(\theta \mathbf{X}) \leq \theta^{-1} \log \mathbb{E} e^{\lambda_{\max}(\theta \mathbf{X})} \leq \theta^{-1} \log \mathbb{E} \operatorname{tr} e^{\theta \mathbf{X}}.$$

The justification is the same as above. Identify the normalized trace mgf, and take the infimum over $\theta > 0$. Similar considerations yield (3.4). \square

3.3. Studying the Trace Mgf with Exchangeable Pairs. The technical difficulty in the matrix Laplace transform method arises because we need to estimate the trace mgf. Previous authors have applied deep results from matrix analysis to accomplish this bound: the Golden-Thompson inequality is central to [AW02, Oli09, Oli10], while Lieb's theorem [Lie73, Thm. 6] animates [Tro11a, Tro11b, HKZ11].

In this paper, we develop a fundamentally different technique for studying the trace mgf. The main idea is to control the *growth* of the trace mgf by bounding its *derivative*. To see why we have adopted this strategy, consider a random Hermitian matrix \mathbf{X} , and observe that the derivative of its trace mgf can be written as

$$m'(\theta) = \mathbb{E} \bar{\operatorname{tr}} [\mathbf{X} e^{\theta \mathbf{X}}]$$

under appropriate regularity conditions. This expression has just the form that we need to invoke the method of exchangeable pairs, Lemma 2.3, with $\mathbf{F}(\mathbf{X}) = e^{\theta \mathbf{X}}$. We obtain

$$m'(\theta) = \frac{1}{2\alpha} \mathbb{E} \bar{\operatorname{tr}} [(\mathbf{X} - \mathbf{X}') (e^{\theta \mathbf{X}} - e^{\theta \mathbf{X}'})]. \quad (3.5)$$

This formula strongly suggests that we should apply a mean value theorem to control the derivative; we establish the result that we need in Section 3.4 below. Ultimately, this argument leads to a differential inequality for $m'(\theta)$, which we can integrate to obtain an estimate for $m(\theta)$.

The technique of bounding the derivative of an mgf lies at the heart of the log-Sobolev method for studying concentration phenomena [Led01, Ch. 5]. Recently, Chatterjee [Cha07, Cha08] demonstrated that the method of exchangeable pairs provides another way to control the derivative of an mgf. Our arguments closely follow the pattern set by Chatterjee; the novelty inheres in the extension of these ideas to the matrix setting and the striking applications that this extension permits.

3.4. The Mean Value Trace Inequality. To bound the expression (3.5) for the derivative of the trace mgf, we need a matrix generalization of the mean value theorem for a function whose derivative is convex. We state the result in full generality because it plays a role later.

Lemma 3.4 (Mean Value Trace Inequality). *Let I be an interval of the real line. Suppose that $g : I \rightarrow \mathbb{R}$ is a weakly increasing function and that $h : I \rightarrow \mathbb{R}$ is a function whose derivative h' is convex. For all matrices $\mathbf{A}, \mathbf{B} \in \mathbb{H}^d(I)$, it holds that*

$$\bar{\text{tr}}[(g(\mathbf{A}) - g(\mathbf{B})) \cdot (h(\mathbf{A}) - h(\mathbf{B}))] \leq \frac{1}{2} \bar{\text{tr}}[(g(\mathbf{A}) - g(\mathbf{B})) \cdot (\mathbf{A} - \mathbf{B}) \cdot (h'(\mathbf{A}) + h'(\mathbf{B}))].$$

When h' is concave, the inequality is reversed. The same results hold for the standard trace.

To prove Lemma 3.4, we need a standard trace inequality [Pet94, Prop. 3], which is an easy consequence of the spectral theorem for Hermitian matrices.

Proposition 3.5 (Generalized Klein Inequality). *Let u_1, \dots, u_n and v_1, \dots, v_n be real-valued functions on an interval I of the real line. Suppose that*

$$\sum_k u_k(a) v_k(b) \geq 0 \quad \text{for all } a, b \in I. \quad (3.6)$$

Then

$$\bar{\text{tr}}\left[\sum_k u_k(\mathbf{A}) v_k(\mathbf{B})\right] \geq 0 \quad \text{for all } \mathbf{A}, \mathbf{B} \in \mathbb{H}^d(I).$$

With the generalized Klein inequality at hand, we can establish Lemma 3.4 by developing the appropriate scalar inequality.

Proof of Lemma 3.4. Fix $a, b \in I$. Since g is weakly increasing, $(g(a) - g(b)) \cdot (a - b) \geq 0$. The fundamental theorem of calculus and the convexity of h' yield the estimate

$$\begin{aligned} (g(a) - g(b)) \cdot (h(a) - h(b)) &= (g(a) - g(b)) \cdot (a - b) \int_0^1 h'(\tau a + (1 - \tau)b) d\tau \\ &\leq (g(a) - g(b)) \cdot (a - b) \int_0^1 [\tau \cdot h'(a) + (1 - \tau) \cdot h'(b)] d\tau \\ &= \frac{1}{2} [(g(a) - g(b)) \cdot (a - b) \cdot (h'(a) + h'(b))]. \end{aligned} \quad (3.7)$$

The inequality is reversed when h' is concave.

The bound (3.7) can be written in the form (3.6) by expanding the products and collecting terms depending on a into functions $u_k(a)$ and terms depending on b into functions $v_k(b)$. Proposition 3.5 then delivers a trace inequality, which can be massaged into the desired form using the cyclicity of the trace and the fact that standard functions of the same matrix commute. We omit the algebraic details. \square

Remark 3.6. We must warn the reader that the proof of Lemma 3.4 succeeds because the trace contains a product of *three* terms involving *two* matrices. The obstacle to proving more general results is that we cannot reorganize expressions like $\text{tr}(\mathbf{A}\mathbf{B}\mathbf{A}\mathbf{B})$ and $\text{tr}(\mathbf{A}\mathbf{B}\mathbf{C})$ at will.

3.5. Bounding the Derivative of the Trace Mgf. The central result in this section applies the method of exchangeable pairs and the mean value trace inequality to bound the derivative of the trace mgf in terms of the conditional variance. This is the most important step in our theory on the exponential concentration of random matrices.

Lemma 3.7 (The Derivative of the Trace Mgf). *Suppose that $(\mathbf{X}, \mathbf{X}') \in \mathbb{H}^d \times \mathbb{H}^d$ is a matrix Stein pair, and assume that \mathbf{X} is almost surely bounded in norm. Define the normalized trace mgf $m(\theta) := \mathbb{E} \bar{\text{tr}} e^{\theta \mathbf{X}}$. Then*

$$m'(\theta) \leq \theta \cdot \mathbb{E} \bar{\text{tr}} [\Delta_{\mathbf{X}} e^{\theta \mathbf{X}}] \quad \text{when } \theta \geq 0. \quad (3.8)$$

$$m'(\theta) \geq \theta \cdot \mathbb{E} \bar{\text{tr}} [\Delta_{\mathbf{X}} e^{\theta \mathbf{X}}] \quad \text{when } \theta \leq 0. \quad (3.9)$$

The conditional variance $\Delta_{\mathbf{X}}$ is defined in (2.4).

Proof. Let us begin with the expression for the derivative of the trace mgf:

$$m'(\theta) = \mathbb{E} \bar{\text{tr}} \left[\frac{d}{d\theta} e^{\theta \mathbf{X}} \right] = \mathbb{E} \bar{\text{tr}} [\mathbf{X} e^{\theta \mathbf{X}}]. \quad (3.10)$$

We can move the derivative inside the expectation because of the dominated convergence theorem and the boundedness of \mathbf{X} .

Now, we apply the method of exchangeable pairs, Lemma 2.3, with $\mathbf{F}(\mathbf{X}) = e^{\theta \mathbf{X}}$ to identify an alternative representation of the derivative (3.10):

$$m'(\theta) = \frac{1}{2\alpha} \mathbb{E} \bar{\text{tr}} [(\mathbf{X} - \mathbf{X}') (e^{\theta \mathbf{X}} - e^{\theta \mathbf{X}'})]. \quad (3.11)$$

We have used the boundedness of \mathbf{X} to verify the regularity condition (2.2).

The expression (3.11) is perfectly suited for an application of the mean value trace inequality, Lemma 3.4. First, assume that $\theta \geq 0$, and consider the function $h : s \mapsto e^{\theta s}$. The derivative $h' : s \mapsto \theta e^{\theta s}$ is convex, so Lemma 3.4 implies that

$$\begin{aligned} m'(\theta) &\leq \frac{\theta}{4\alpha} \mathbb{E} \bar{\text{tr}} [(\mathbf{X} - \mathbf{X}')^2 \cdot (e^{\theta \mathbf{X}} + e^{\theta \mathbf{X}'})] \\ &= \frac{\theta}{2\alpha} \mathbb{E} \bar{\text{tr}} [(\mathbf{X} - \mathbf{X}')^2 \cdot e^{\theta \mathbf{X}}] \\ &= \theta \cdot \mathbb{E} \bar{\text{tr}} \left[\frac{1}{2\alpha} \mathbb{E} [(\mathbf{X} - \mathbf{X}')^2 | Z] \cdot e^{\theta \mathbf{X}} \right]. \end{aligned}$$

The second line follows from the fact that $(\mathbf{X}, \mathbf{X}')$ is an exchangeable pair. In the last line, we have used the boundedness of \mathbf{X} and \mathbf{X}' to invoke the pull-through property of conditional expectation. Identify the conditional variance $\Delta_{\mathbf{X}}$, defined in (2.4), to complete the argument.

The result for $\theta \leq 0$ follows from an analogous argument. In this case, we simply observe that the derivative of the function $h : s \mapsto e^{\theta s}$ is now concave, so the mean value trace inequality, Lemma 3.4, produces a lower bound. The remaining steps are identical. \square

Remark 3.8 (Regularity Conditions). To simplify the presentation, we have instated a boundedness assumption in Lemma 3.7. All the examples we discuss satisfy this requirement. When \mathbf{X} is unbounded, Lemma 3.7 still holds provided that \mathbf{X} meets an integrability condition.

4. EXPONENTIAL CONCENTRATION FOR BOUNDED RANDOM MATRICES

We are now prepared to establish exponential concentration inequalities. Our first major result demonstrates that an almost-sure bound for the conditional variance yields exponential tail bounds for the extreme eigenvalues of a random Hermitian matrix. We can also obtain estimates for the expectation of the extreme eigenvalues.

Theorem 4.1 (Concentration for Bounded Random Matrices). *Consider a matrix Stein pair $(\mathbf{X}, \mathbf{X}') \in \mathbb{H}^d \times \mathbb{H}^d$. Suppose there exist nonnegative constants c, v for which the conditional variance (2.4) of the pair satisfies*

$$\Delta_{\mathbf{X}} \preceq c\mathbf{X} + v\mathbf{I} \quad \text{almost surely.} \quad (4.1)$$

Then, for all $t \geq 0$,

$$\begin{aligned} \mathbb{P}\{\lambda_{\min}(\mathbf{X}) \leq -t\} &\leq d \cdot \exp\left\{\frac{-t^2}{2v}\right\}. \\ \mathbb{P}\{\lambda_{\max}(\mathbf{X}) \geq t\} &\leq d \cdot \exp\left\{-\frac{t}{c} + \frac{v}{c^2} \log\left(1 + \frac{ct}{v}\right)\right\} \\ &\leq d \cdot \exp\left\{\frac{-t^2}{2v + 2ct}\right\}. \end{aligned}$$

Furthermore,

$$\begin{aligned} \mathbb{E} \lambda_{\min}(\mathbf{X}) &\geq -\sqrt{2v \log d} \\ \mathbb{E} \lambda_{\max}(\mathbf{X}) &\leq \sqrt{2v \log d} + c \log d. \end{aligned}$$

This result may be viewed as a matrix analogue of Chatterjee's concentration inequality for scalar random variables [Cha07, Thm. 1.5(ii)]. The proof of Theorem 4.1 appears below in Section 4.2. Before we present the argument, let us explain how the result provides a short proof of a Hoeffding-type inequality for matrices.

4.1. Application: Matrix Hoeffding Inequality. Theorem 4.1 yields an extension of Hoeffding's inequality [Hoe63] that holds for an independent sum of bounded random matrices.

Corollary 4.2 (Matrix Hoeffding). *Consider a finite sequence $(\mathbf{Y}_k)_{k \geq 1}$ of independent random matrices in \mathbb{H}^d and a finite sequence $(\mathbf{A}_k)_{k \geq 1}$ of deterministic matrices in \mathbb{H}^d . Assume that*

$$\mathbb{E} \mathbf{Y}_k = \mathbf{0} \quad \text{and} \quad \mathbf{Y}_k^2 \preceq \mathbf{A}_k^2 \quad \text{for each index } k.$$

Then, for all $t \geq 0$,

$$\mathbb{P}\left\{\lambda_{\max}\left(\sum_k \mathbf{Y}_k\right) \geq t\right\} \leq d \cdot e^{-t^2/2\sigma^2} \quad \text{where} \quad \sigma^2 := \frac{1}{2} \left\|\sum_k (\mathbf{A}_k^2 + \mathbb{E} \mathbf{Y}_k^2)\right\|.$$

Furthermore,

$$\mathbb{E} \lambda_{\max}\left(\sum_k \mathbf{Y}_k\right) \leq \sigma \sqrt{2 \log d}.$$

Proof. Let $\mathbf{X} = \sum_k \mathbf{Y}_k$. Since \mathbf{X} is a sum of centered, independent random matrices, we can use the matrix Stein pair constructed in Section 2.4. According to (2.6), the conditional variance satisfies

$$\Delta_{\mathbf{X}} = \frac{1}{2} \sum_k (\mathbf{Y}_k^2 + \mathbb{E} \mathbf{Y}_k^2) \preceq \sigma^2 \mathbf{I}$$

because $\mathbf{Y}_k^2 \preceq \mathbf{A}_k^2$. Invoke Theorem 4.1 with $c = 0$ and $v = \sigma^2$ to complete the bound. \square

In the scalar setting $d = 1$, Corollary 4.2 reproduces to an inequality of Chatterjee [Cha07, Sec. 1.5], which itself represents an improvement over the classical scalar Hoeffding bound. In turn, Corollary 4.2 improves upon the matrix Hoeffding inequality of [Tro11b, Thm. 1.3] in two ways. First, we have reduced the constant in the exponent to its optimal value $1/2$. Second, we have decreased the size of the variance measure because $\sigma^2 \leq \left\|\sum_k \mathbf{A}_k^2\right\|$.

4.2. Proof of Theorem 4.1: Exponential Concentration. Suppose that $(\mathbf{X}, \mathbf{X}')$ is a matrix Stein pair constructed from an auxiliary exchangeable pair (Z, Z') . Our aim is to bound the normalized trace mgf

$$m(\theta) := \mathbb{E} \bar{\text{tr}} e^{\theta \mathbf{X}} \quad \text{for } \theta \in \mathbb{R}. \quad (4.2)$$

The basic strategy is to develop a differential inequality, which we integrate to control $m(\theta)$ itself. Once these estimates are in place, the matrix Laplace transform method, Proposition 3.3, furnishes probability inequalities for the extreme eigenvalues of \mathbf{X} .

The following result summarizes our bounds for the trace mgf $m(\theta)$.

Lemma 4.3 (Trace Mgf Estimates for Bounded Random Matrices). *Let $(\mathbf{X}, \mathbf{X}')$ be a matrix Stein pair, and suppose there exist nonnegative constants c, v for which*

$$\Delta_{\mathbf{X}} \preceq c\mathbf{X} + v\mathbf{I} \quad \text{almost surely.} \quad (4.3)$$

Then the normalized trace mgf $m(\theta) := \mathbb{E} \bar{\text{tr}} e^{\theta \mathbf{X}}$ satisfies the bounds

$$\log m(\theta) \leq \frac{v\theta^2}{2} \quad \text{when } \theta \leq 0. \quad (4.4)$$

$$\log m(\theta) \leq \frac{v}{c^2} \left[\log \left(\frac{1}{1 - c\theta} \right) - c\theta \right] \quad (4.5)$$

$$\leq \frac{v\theta^2}{2(1 - c\theta)} \quad \text{when } 0 \leq \theta < 1/c. \quad (4.6)$$

We establish Lemma 4.3 in Section 4.2.1 et seq. In Section 4.2.4, we finish the proof of Theorem 4.1 by combining these bounds with the matrix Laplace transform method.

4.2.1. Boundedness of the Random Matrix. First, we confirm that the random matrix \mathbf{X} is almost surely bounded under the hypothesis (4.3) on the conditional variance $\Delta_{\mathbf{X}}$. Recall the definition (2.4) of the conditional variance, and compute that

$$\Delta_{\mathbf{X}} = \frac{1}{2\alpha} \mathbb{E}[(\mathbf{X} - \mathbf{X}')^2 | Z] \succcurlyeq \frac{1}{2\alpha} (\mathbb{E}[\mathbf{X} - \mathbf{X}' | Z])^2 = \frac{1}{2\alpha} (\alpha \mathbf{X})^2 = \frac{\alpha}{2} \mathbf{X}^2.$$

The semidefinite bound is the operator Jensen inequality (1.2), applied conditionally. The third relation follows from the definition (2.1) of a matrix Stein pair. Owing to the assumption (4.3), we reach the quadratic inequality $\frac{1}{2}\alpha \mathbf{X}^2 \preceq c\mathbf{X} + v\mathbf{I}$. The scale factor α is positive, so we may conclude that the eigenvalues of \mathbf{X} are almost surely restricted to a bounded interval.

4.2.2. Differential Inequalities for the Trace Mgf. The fact that \mathbf{X} is almost surely bounded ensures that the derivative of the trace mgf has the form

$$m'(\theta) = \mathbb{E} \bar{\text{tr}} [\mathbf{X} e^{\theta \mathbf{X}}] \quad \text{for } \theta \in \mathbb{R}. \quad (4.7)$$

To bound the derivative, we combine Lemma 3.7 with the assumed inequality (4.3) for the conditional variance. For $\theta \geq 0$, we obtain

$$\begin{aligned} m'(\theta) &\leq \theta \cdot \mathbb{E} \bar{\text{tr}} [\Delta_{\mathbf{X}} e^{\theta \mathbf{X}}] \\ &\leq \theta \cdot \mathbb{E} \bar{\text{tr}} [(c\mathbf{X} + v\mathbf{I}) e^{\theta \mathbf{X}}] \\ &= c\theta \cdot \mathbb{E} \bar{\text{tr}} [\mathbf{X} e^{\theta \mathbf{X}}] + v\theta \cdot \mathbb{E} \bar{\text{tr}} e^{\theta \mathbf{X}} \\ &= c\theta \cdot m'(\theta) + v\theta \cdot m(\theta). \end{aligned}$$

The second relation relies on the fact that the matrix $e^{\theta \mathbf{X}}$ is positive. In the last line, we have identified the trace mgf (4.2) and its derivative (4.7). For $\theta \leq 0$, the same argument yields a lower bound

$$m'(\theta) \geq c\theta \cdot m'(\theta) + v\theta \cdot m(\theta).$$

Rearrange these inequalities to isolate the log-derivative $m'(\theta)/m(\theta)$ of the trace mgf. We reach

$$\frac{d}{d\theta} \log m(\theta) \leq \frac{v\theta}{1-c\theta} \quad \text{for } 0 \leq \theta < 1/c, \text{ and} \quad (4.8)$$

$$\frac{d}{d\theta} \log m(\theta) \geq \frac{v\theta}{1-c\theta} \quad \text{for } \theta \leq 0. \quad (4.9)$$

4.2.3. *Solving the Differential Inequalities.* Observe that

$$\log m(0) = \log \bar{\text{tr}} \mathbf{e}^0 = \log \bar{\text{tr}} \mathbf{I} = \log 1 = 0. \quad (4.10)$$

Therefore, we may integrate the differential inequalities (4.8) and (4.9), starting at zero, to obtain bounds on $\log m(\theta)$ elsewhere.

First, assume that $0 \leq \theta < 1/c$. In view of (4.10), the fundamental theorem of calculus and the differential inequality (4.8) imply that

$$\log m(\theta) = \int_0^\theta \frac{d}{ds} \log m(s) ds \leq \int_0^\theta \frac{vs}{1-cs} ds = -\frac{v}{c^2}(c\theta + \log(1-c\theta)).$$

We can develop a weaker inequality by making a further approximation within the integral:

$$\log m(\theta) \leq \int_0^\theta \frac{vs}{1-cs} ds \leq \int_0^\theta \frac{vs}{1-c\theta} ds \leq \frac{v\theta^2}{2(1-c\theta)}.$$

These inequalities are the trace mgf estimates (4.5) and (4.6) appearing in Lemma 4.3.

Next, assume that $\theta \leq 0$. In this case, the differential inequality (4.9) yields

$$-\log m(\theta) = \int_\theta^0 \frac{d}{ds} \log m(s) ds \geq \int_\theta^0 \frac{vs}{1-cs} ds \geq \int_\theta^0 vs ds = -\frac{v\theta^2}{2}.$$

This calculation delivers the trace mgf bound (4.4), and the proof of Lemma 4.3 is complete.

4.2.4. *The Matrix Laplace Transform Argument.* With Lemma 4.3 at hand, we quickly finish the proof of Theorem 4.1. First, let us establish probability inequalities for the maximum eigenvalue. The Laplace transform bound (3.1) and the trace mgf estimate (4.5) together yield

$$\mathbb{P}\{\lambda_{\max}(\mathbf{X}) \geq t\} \leq d \cdot \inf_{\theta > 0} \exp\left\{-\theta t - \frac{v}{c^2}(c\theta + \log(1-c\theta))\right\} \leq d \cdot \exp\left\{-\frac{t}{c} + \frac{v}{c^2} \log\left(1 + \frac{ct}{v}\right)\right\}.$$

The second relation follows when we choose $\theta = t/(v+ct)$. Similarly, the trace mgf bound (4.6) delivers

$$\mathbb{P}\{\lambda_{\max}(\mathbf{X}) \geq t\} \leq d \cdot \inf_{\theta > 0} \exp\left\{-\theta t + \frac{v\theta^2}{2(1-c\theta)}\right\} \leq d \cdot \exp\left\{-\frac{t^2}{2v+2ct}\right\},$$

where we have selected $\theta = t/(v+ct)$ again. To control the expectation of the maximum eigenvalue, we invoke the Laplace transform bound (3.3) and the trace mgf bound (4.6) to see that

$$\mathbb{E} \lambda_{\max}(\mathbf{X}) \leq \inf_{\theta > 0} \frac{1}{\theta} \left[\log d + \frac{v\theta^2}{2(1-c\theta)} \right] = \sqrt{2v \log d} + c \log d.$$

The second relation can be verified using a computer algebra system.

Next, we turn to the results for the minimum eigenvalue. Combine the matrix Laplace transform bound (3.2) with the trace mgf bound (4.4) to reach

$$\mathbb{P}\{\lambda_{\min}(\mathbf{X}) \leq -t\} \leq d \cdot \inf_{\theta < 0} \exp\left\{\theta t + \frac{v\theta^2}{2}\right\} = d \cdot e^{-t^2/2v}.$$

The infimum is attained at $\theta = -t/v$. To compute the expectation of the minimum eigenvalue, we apply the Laplace transform bound (3.4) and the trace mgf bound (4.4), whence

$$\mathbb{E} \lambda_{\min}(\mathbf{X}) \geq \sup_{\theta < 0} \frac{1}{\theta} \left[\log d + \frac{v\theta^2}{2} \right] = -\sqrt{2v \log d}.$$

The infimum is attained at $\theta = -\sqrt{2v^{-1} \log d}$.

5. REFINED EXPONENTIAL CONCENTRATION FOR BOUNDED RANDOM MATRICES

Although Theorem 4.1 is a strong result, the hypothesis $\Delta_{\mathbf{X}} \preceq c\mathbf{X} + v\mathbf{I}$ on the conditional variance is too stringent for many situations of interest. Our second major result shows that we can use the typical behavior of the conditional variance to obtain tail bounds for the maximum eigenvalue of a random Hermitian matrix.

Theorem 5.1 (Refined Concentration for Random Matrices). *Let $(\mathbf{X}, \mathbf{X}') \in \mathbb{H}^d \times \mathbb{H}^d$ be a matrix Stein pair, and assume that \mathbf{X} is almost surely bounded in norm. Define the function*

$$r(\psi) := \frac{1}{\psi} \log \mathbb{E} \operatorname{tr} e^{\psi \Delta_{\mathbf{X}}} \quad \text{for each } \psi > 0, \quad (5.1)$$

where $\Delta_{\mathbf{X}}$ is the conditional variance (2.4). Then, for all $t \geq 0$ and all $\psi > 0$,

$$\mathbb{P} \{ \lambda_{\max}(\mathbf{X}) \geq t \} \leq d \cdot \exp \left\{ \frac{-t^2}{2r(\psi) + 2t/\sqrt{\psi}} \right\}. \quad (5.2)$$

Furthermore, for all $\psi > 0$,

$$\mathbb{E} \lambda_{\max}(\mathbf{X}) \leq \sqrt{2r(\psi) \log d} + \frac{\log d}{\sqrt{\psi}}. \quad (5.3)$$

This theorem is essentially a matrix version of a result from Chatterjee's thesis [Cha08, Thm. 3.13]. The proof of Theorem 5.1 is similar in spirit to the proof of Theorem 4.1, so we postpone the demonstration until Appendix A.

Let us offer some remarks to clarify the meaning of this result. Recall that $\Delta_{\mathbf{X}}$ is a stochastic approximation for the variance of the random matrix \mathbf{X} . We can interpret the function $r(\psi)$ as a measure of the typical magnitude of the conditional variance. Indeed, the matrix Laplace transform result, Proposition 3.3, ensures that

$$\mathbb{E} \lambda_{\max}(\Delta_{\mathbf{X}}) \leq \inf_{\psi > 0} \left[r(\psi) + \frac{\log d}{\psi} \right].$$

The moral of this inequality is that we can often identify a value of ψ to make $r(\psi) \approx \mathbb{E} \lambda_{\max}(\Delta_{\mathbf{X}})$. Ideally, we also want to choose $r(\psi) \gg \psi^{-1/2}$ so that the term $r(\psi)$ drives the tail bound (5.2) when the parameter t is small. In the next subsection, we show that these heuristics yield a matrix Bernstein inequality.

5.1. Application: The Matrix Bernstein Inequality. As an illustration of Theorem 5.1, we establish a tail bound for a sum of centered, independent random matrices that are subject to a uniform norm bound.

Corollary 5.2 (Matrix Bernstein). *Consider an independent sequence $(\mathbf{Y}_k)_{k \geq 1}$ of random matrices in \mathbb{H}^d that satisfy*

$$\mathbb{E} \mathbf{Y}_k = \mathbf{0} \quad \text{and} \quad \|\mathbf{Y}_k\| \leq R \quad \text{for each index } k.$$

Then, for all $t \geq 0$,

$$\mathbb{P} \left\{ \lambda_{\max} \left(\sum_k \mathbf{Y}_k \right) \geq t \right\} \leq d \cdot \exp \left\{ \frac{-t^2}{3\sigma^2 + 2Rt} \right\} \quad \text{where} \quad \sigma^2 := \left\| \sum_k \mathbb{E} \mathbf{Y}_k^2 \right\|.$$

Furthermore,

$$\mathbb{E} \lambda_{\max} \left(\sum_k \mathbf{Y}_k \right) \leq \sigma \sqrt{3 \log d} + R \log d.$$

Corollary 5.2 is directly comparable with other matrix Bernstein inequalities in the literature. The constants here are slightly worse than [Tro11b, Thm. 1.4] and slightly better than [Oli09, Thm. 1.2]. The hypotheses in the current result are somewhat stricter than those in the prior works. Nevertheless, the proof provides a template for studying more complicated random matrices that involve dependent random variables.

Proof. We consider the matrix Stein pair $(\mathbf{X}, \mathbf{X}')$ described in Section 2.4. The calculation (2.6) shows that the conditional variance of \mathbf{X} satisfies

$$\Delta_{\mathbf{X}} = \frac{1}{2} \sum_k (\mathbf{Y}_k^2 + \mathbb{E} \mathbf{Y}_k^2).$$

The function $r(\psi)$ measures the typical size of $\Delta_{\mathbf{X}}$. To control $r(\psi)$, we center the conditional variance and reduce the expression as follows.

$$\begin{aligned} r(\psi) &:= \frac{1}{\psi} \log \mathbb{E} \bar{\text{tr}} e^{\psi \Delta_{\mathbf{X}}} \leq \frac{1}{\psi} \log \mathbb{E} \bar{\text{tr}} \exp \{ \psi (\Delta_{\mathbf{X}} - \mathbb{E} \Delta_{\mathbf{X}}) + \psi \|\mathbb{E} \Delta_{\mathbf{X}}\| \cdot \mathbf{I} \} \\ &= \frac{1}{\psi} \log \mathbb{E} \bar{\text{tr}} [e^{\psi \sigma^2} \cdot \exp \{ \psi (\Delta_{\mathbf{X}} - \mathbb{E} \Delta_{\mathbf{X}}) \}] \\ &= \sigma^2 + \frac{1}{\psi} \log \mathbb{E} \bar{\text{tr}} e^{\psi (\Delta_{\mathbf{X}} - \mathbb{E} \Delta_{\mathbf{X}})}. \end{aligned} \quad (5.4)$$

The inequality depends on the monotonicity of the trace exponential [Pet94, Sec. 2]. Afterward, we have applied the identity $\|\mathbb{E} \Delta_{\mathbf{X}}\| = \|\mathbb{E} \mathbf{X}^2\| = \sigma^2$, which follows from (2.5) and the independence of the sequence $(\mathbf{Y}_k)_{k \geq 1}$.

Introduce the centered random matrix

$$\mathbf{W} := \Delta_{\mathbf{X}} - \mathbb{E} \Delta_{\mathbf{X}} = \frac{1}{2} \sum_k (\mathbf{Y}_k^2 - \mathbb{E} \mathbf{Y}_k^2). \quad (5.5)$$

Observe that \mathbf{W} consists of a sum of centered, independent random matrices, so we can study it using the matrix Stein pair discussed in Section 2.4. Adapt the conditional variance calculation (2.6) to obtain

$$\begin{aligned} \Delta_{\mathbf{W}} &= \frac{1}{2} \cdot \frac{1}{4} \sum_k \left[(\mathbf{Y}_k^2 - \mathbb{E} \mathbf{Y}_k^2)^2 + \mathbb{E} (\mathbf{Y}_k^2 - \mathbb{E} \mathbf{Y}_k^2)^2 \right] \\ &\preceq \frac{1}{8} \sum_k \left[2\mathbf{Y}_k^4 + 2(\mathbb{E} \mathbf{Y}_k^2)^2 + \mathbb{E} \mathbf{Y}_k^4 - (\mathbb{E} \mathbf{Y}_k^2)^2 \right] \\ &\preceq \frac{1}{4} \sum_k (\mathbf{Y}_k^4 + \mathbb{E} \mathbf{Y}_k^4). \end{aligned}$$

To reach the second line, we apply the operator convexity (1.1) of the matrix square to the first parenthesis, and we compute the second expectation explicitly. The third line follows from the operator Jensen inequality (1.2). To continue, make the estimate $\mathbf{Y}_k^4 \preceq R^2 \mathbf{Y}_k^2$ in both terms. Thus,

$$\Delta_{\mathbf{W}} \preceq \frac{R^2}{4} \sum_{k=1}^n (\mathbf{Y}_k^2 + \mathbb{E} \mathbf{Y}_k^2) \preceq \frac{R^2}{2} \cdot \mathbf{W} + \frac{R^2 \sigma^2}{2} \cdot \mathbf{I}.$$

The trace mgf bound, Lemma 4.3, delivers

$$\log m_{\mathbf{W}}(\psi) = \log \mathbb{E} \bar{\text{tr}} e^{\psi \mathbf{W}} \leq \frac{R^2 \sigma^2 \psi^2}{4 - 2R^2 \psi}. \quad (5.6)$$

To complete the proof, combine the bounds (5.4) and (5.6) to reach

$$r(\psi) \leq \sigma^2 + \frac{R^2 \sigma^2 \psi}{4 - 2R^2 \psi}.$$

In particular, it holds that $r(R^{-2}) \leq 1.5\sigma^2$. The result now follows from Theorem 5.1. \square

6. POLYNOMIAL MOMENTS AND THE SPECTRAL NORM OF A RANDOM MATRIX

We can also study the spectral norm of a random matrix by bounding its polynomial moments. To progress toward these results, let us introduce the family of Schatten norms.

Definition 6.1 (Schatten Norm). For each $p \geq 1$, the Schatten p -norm is defined as

$$\|\mathbf{B}\|_p := (\operatorname{tr} |\mathbf{B}|^p)^{1/p} \quad \text{for } \mathbf{B} \in \mathbb{M}^d.$$

In this setting, $|\mathbf{B}| := (\mathbf{B}^* \mathbf{B})^{1/2}$. Bhatia's book [Bha97, Ch. IV] contains a detailed discussion of these norms and their properties.

The following proposition is a matrix analog of the Chebyshev bound from classical probability.

Proposition 6.2 (Matrix Chebyshev Method). *Let \mathbf{X} be a random matrix. For all $t > 0$,*

$$\mathbb{P}\{\|\mathbf{X}\| \geq t\} \leq \inf_{p \geq 1} t^{-p} \cdot \mathbb{E} \|\mathbf{X}\|_p^p. \quad (6.1)$$

Furthermore,

$$\mathbb{E} \|\mathbf{X}\| \leq \inf_{p \geq 1} (\mathbb{E} \|\mathbf{X}\|_p^p)^{1/p}. \quad (6.2)$$

Proof. To prove (6.1), we use Markov's inequality. For $p \geq 1$,

$$\mathbb{P}\{\|\mathbf{X}\| \geq t\} \leq t^{-p} \cdot \mathbb{E} \|\mathbf{X}\|^p = t^{-p} \cdot \mathbb{E} \|\mathbf{X}^p\| \leq t^{-p} \cdot \mathbb{E} \operatorname{tr} |\mathbf{X}|^p,$$

since the trace of a positive matrix dominates the maximum eigenvalue. To verify (6.2), select $p \geq 1$. Jensen's inequality implies that

$$\mathbb{E} \|\mathbf{X}\| \leq (\mathbb{E} \|\mathbf{X}\|^p)^{1/p} = (\mathbb{E} \|\mathbf{X}^p\|)^{1/p} \leq (\mathbb{E} \operatorname{tr} |\mathbf{X}|^p)^{1/p}.$$

Identify the Schatten p -norm and take infima to complete the bounds. \square

Remark 6.3 (Chebyshev vs. Laplace). The matrix Chebyshev bound (6.1) is at least as tight as the analogous matrix Laplace transform bound (3.1) from Proposition 3.3. Indeed, suppose that \mathbf{X} is a bounded random matrix. For all $\theta, t > 0$, we have

$$\begin{aligned} e^{-\theta t} \cdot \mathbb{E} \operatorname{tr} e^{\theta |\mathbf{X}|} &= e^{-\theta t} \cdot \sum_{q=0}^{\infty} \frac{\theta^q}{q!} (t^q t^{-q}) \mathbb{E} \operatorname{tr} |\mathbf{X}|^q \\ &\geq e^{-\theta t} \cdot \sum_{q=0}^{\infty} \frac{(\theta t)^q}{q!} \cdot \left[1 \wedge \inf_{p \geq 1} (t^{-p} \cdot \mathbb{E} \operatorname{tr} |\mathbf{X}|^p) \right] = 1 \wedge \inf_{p \geq 1} (t^{-p} \cdot \mathbb{E} \|\mathbf{X}\|_p^p). \end{aligned}$$

The first identity follows from the Taylor expansion of the matrix exponential. A similar argument allows us to convert polynomial moment bounds into bounds on the trace mgf.

7. POLYNOMIAL MOMENT INEQUALITIES FOR RANDOM MATRICES

Our last major result demonstrates that the polynomial moments of a random Hermitian matrix are controlled by the moments of the conditional variance. By combining this result with the matrix Chebyshev method, Proposition 6.2, we can obtain probability inequalities for the spectral norm of a random Hermitian matrix.

Theorem 7.1 (Matrix BDG Inequality). *Let $p = 1$ or $p \geq 1.5$. Suppose that $(\mathbf{X}, \mathbf{X}')$ is a matrix Stein pair where $\mathbb{E} \|\mathbf{X}\|_{2p}^{2p} < \infty$. Then*

$$(\mathbb{E} \|\mathbf{X}\|_{2p}^{2p})^{1/2p} \leq \sqrt{2p-1} \cdot (\mathbb{E} \|\Delta_{\mathbf{X}}\|_p^p)^{1/2p}.$$

The conditional variance $\Delta_{\mathbf{X}}$ is defined in (2.4).

This theorem extends a result of Chatterjee [Cha07, Thm. 1.5(iii)] to the matrix setting. Chatterjee's bound can be viewed as an exchangeable pairs version of the Burkholder–Davis–Gundy (BDG) inequality from classical martingale theory [Bur73]. Matrix extensions of the BDG inequality appear in the work of Pisier–Xu [PX97] and the work of Junge–Xu [JX03, JX08].

The proof of Theorem 7.1 appears below in Section 7.3. Before we present the argument, let us offer a few tangential remarks and describe a few striking applications of this inequality.

Remark 7.2 (Missing Values). The Matrix BDG inequality also holds when $1 < p < 1.5$. In this range, our best bound for the constant is $\sqrt{4p-2}$. The proof requires a variant of the mean value trace inequality for a convex function h .

Remark 7.3 (Infinite-Dimensional Versions). Theorem 7.1 and its corollaries extend to Schatten-class operators.

7.1. Application: Matrix Khintchine Inequality. First, we demonstrate that the matrix BDG inequality contains an improvement of the noncommutative Khintchine inequality [LP86, LPP91] in the matrix setting. This result has been a dominant tool in several application areas over the last few years, largely because of the articles [Rud99, RV07].

Corollary 7.4 (Matrix Khintchine). *Suppose that $p = 1$ or $p \geq 1.5$. Consider a finite sequence $(\mathbf{Y}_k)_{k \geq 1}$ of independent, random, Hermitian matrices and a deterministic sequence $(\mathbf{A}_k)_{k \geq 1}$ for which*

$$\mathbb{E} \mathbf{Y}_k = \mathbf{0} \quad \text{and} \quad \mathbf{Y}_k^2 \preceq \mathbf{A}_k^2 \quad \text{for each index } k. \quad (7.1)$$

Then

$$\left(\mathbb{E} \left\| \sum_k \mathbf{Y}_k \right\|_{2p}^{2p} \right)^{1/2p} \leq \sqrt{p-0.5} \cdot \left\| \left(\sum_k (\mathbf{A}_k^2 + \mathbb{E} \mathbf{Y}_k^2) \right)^{1/2} \right\|_{2p}.$$

In particular, when $(\varepsilon_k)_{k \geq 1}$ is an independent sequence of Rademacher random variables,

$$\left(\mathbb{E} \left\| \sum_k \varepsilon_k \mathbf{A}_k \right\|_{2p}^{2p} \right)^{1/2p} \leq \sqrt{2p-1} \cdot \left\| \left(\sum_k \mathbf{A}_k^2 \right)^{1/2} \right\|_{2p}. \quad (7.2)$$

Proof. Consider the random matrix $\mathbf{X} = \sum_k \mathbf{Y}_k$. We use the matrix Stein pair constructed in Section 2.4. According to (2.6), the conditional variance $\Delta_{\mathbf{X}}$ satisfies

$$\Delta_{\mathbf{X}} = \frac{1}{2} \sum_k (\mathbf{Y}_k^2 + \mathbb{E} \mathbf{Y}_k^2) \preceq \frac{1}{2} \sum_k (\mathbf{A}_k^2 + \mathbb{E} \mathbf{Y}_k^2).$$

An application of Theorem 7.1 completes the argument. \square

Buchholz [Buc01, Thm. 5] has demonstrated that the optimal constant C_{2p} in (7.2) satisfies

$$C_{2p}^{2p} = (2p-1)!! = \frac{(2p)!}{2^p p!} \quad \text{for } p = 1, 2, 3, \dots$$

It can be verified using basic analysis that

$$\frac{(2p-1)^p}{(2p-1)!!} < e^{p-1/2}.$$

As a consequence, the constant in (7.2) lies within a factor \sqrt{e} of optimal. Previous methods for establishing the matrix Khintchine inequality are rather involved, so it is remarkable that the simple argument based on exchangeable pairs leads to a result that is so accurate.

7.2. Application: Matrix Rosenthal Inequality. As a second example, we can develop a more sophisticated set of moment inequalities that are roughly the polynomial equivalent of the exponential moment bound underlying the matrix Bernstein inequality.

Corollary 7.5 (Matrix Rosenthal Inequality). *Suppose that $p = 1$ or $p \geq 1.5$. Consider a finite sequence $(\mathbf{P}_k)_{k \geq 1}$ of independent, random psd matrices that satisfy $\mathbb{E} \|\mathbf{P}_k\|_{2p}^{2p} < \infty$. Then*

$$\left(\mathbb{E} \left\| \sum_k \mathbf{P}_k \right\|_{2p}^{2p} \right)^{1/2p} \leq \left[\left\| \sum_k \mathbb{E} \mathbf{P}_k \right\|_{2p}^{1/2} + \sqrt{4p-2} \cdot \left(\sum_k \mathbb{E} \|\mathbf{P}_k\|_{2p}^{2p} \right)^{1/4p} \right]^2. \quad (7.3)$$

Now, consider a finite sequence $(\mathbf{Y}_k)_{k \geq 1}$ of centered, independent, random Hermitian matrices, and assume that $\mathbb{E} \|\mathbf{Y}_k\|_{4p}^{4p} < \infty$. Then

$$\left(\mathbb{E} \left\| \sum_k \mathbf{Y}_k \right\|_{4p}^{4p} \right)^{1/4p} \leq \sqrt{4p-1} \cdot \left\| \left(\sum_k \mathbb{E} \mathbf{Y}_k^2 \right)^{1/2} \right\|_{4p} + (4p-1) \cdot \left(\sum_k \mathbb{E} \|\mathbf{Y}_k\|_{4p}^{4p} \right)^{1/4p}. \quad (7.4)$$

Turn to Appendix B for our proof of Corollary 7.5. This result extends a moment inequality due to Nagaev and Pinelis [NP77], which refines the constants in Rosenthal's inequality [Ros70, Lem. 1]. See the historical discussion [Pin94, Sec. 5] for details. As we were finishing this paper, we learned that Junge and Zheng have recently established a noncommutative moment inequality [JZ11, Thm. 0.4] that is quite similar to Corollary 7.5.

7.3. Proof of the Matrix BDG Inequality. In many respects, the proof of the matrix BDG inequality is similar to the proof of the exponential concentration result, Theorem 4.1. Both are based on moment comparison arguments that ultimately depend on the method of exchangeable pairs and the mean value trace inequality.

Suppose that $(\mathbf{X}, \mathbf{X}')$ is a matrix Stein pair with scale factor α . First, observe that the result for $p = 1$ already follows from (2.5). Therefore, we may assume that $p \geq 1.5$. Let us introduce notation for the quantity of interest:

$$E := \mathbb{E} \|\mathbf{X}\|_{2p}^{2p} = \mathbb{E} \operatorname{tr} |\mathbf{X}|^{2p}.$$

Rewrite the expression for E by peeling off a copy of $|\mathbf{X}|$. This move yields

$$E = \mathbb{E} \operatorname{tr} [|\mathbf{X}| \cdot |\mathbf{X}|^{2p-1}] = \mathbb{E} \operatorname{tr} [\mathbf{X} \cdot \operatorname{sgn}(\mathbf{X}) \cdot |\mathbf{X}|^{2p-1}].$$

Apply the method of exchangeable pairs, Lemma 2.3, with $\mathbf{F}(\mathbf{X}) = \operatorname{sgn}(\mathbf{X}) \cdot |\mathbf{X}|^{2p-1}$ to reach

$$E = \frac{1}{2\alpha} \mathbb{E} \operatorname{tr} [(\mathbf{X} - \mathbf{X}') \cdot (\operatorname{sgn}(\mathbf{X}) \cdot |\mathbf{X}|^{2p-1} - \operatorname{sgn}(\mathbf{X}') \cdot |\mathbf{X}'|^{2p-1})].$$

To verify the regularity condition (2.2) we need for Lemma 2.3, compute that

$$\begin{aligned} \mathbb{E} \|(\mathbf{X} - \mathbf{X}') \cdot \operatorname{sgn}(\mathbf{X}) \cdot |\mathbf{X}|^{2p-1}\| &\leq \mathbb{E} (\|\mathbf{X}\| \|\mathbf{X}\|^{2p-1}) + \mathbb{E} (\|\mathbf{X}'\| \|\mathbf{X}\|^{2p-1}) \\ &\leq 2(\mathbb{E} \|\mathbf{X}\|^{2p})^{1/2p} (\mathbb{E} \|\mathbf{X}\|^{2p})^{(2p-1)/2p} = 2\mathbb{E} \|\mathbf{X}\|^{2p} < \infty. \end{aligned}$$

We have used the fact that $\operatorname{sgn}(\mathbf{X})$ is a unitary matrix, the exchangeability of $(\mathbf{X}, \mathbf{X}')$, Hölder's inequality for expectation, and the fact that the Schatten $2p$ -norm dominates the spectral norm.

We intend to apply the mean value trace inequality to obtain an estimate for the quantity E . Consider the function $h : s \mapsto \operatorname{sgn}(s) \cdot |s|^{2p-1}$. Its derivative, $h'(s) = (2p-1) \cdot |s|^{2p-2}$, is convex because $p \geq 1.5$. Lemma 3.4 delivers the bound

$$\begin{aligned} E &\leq \frac{2p-1}{4\alpha} \mathbb{E} \operatorname{tr} [(\mathbf{X} - \mathbf{X}')^2 \cdot (|\mathbf{X}|^{2p-2} + |\mathbf{X}'|^{2p-2})] \\ &= \frac{2p-1}{2\alpha} \mathbb{E} \operatorname{tr} [(\mathbf{X} - \mathbf{X}')^2 \cdot |\mathbf{X}|^{2p-2}] \\ &= (2p-1) \cdot \mathbb{E} \operatorname{tr} [\Delta_{\mathbf{X}} \cdot |\mathbf{X}|^{2p-2}]. \end{aligned}$$

The second line follows from the exchangeability of \mathbf{X} and \mathbf{X}' . In the last line, we identify the conditional variance $\Delta_{\mathbf{X}}$, defined in (2.4). As before, the moment bound $\mathbb{E} \|\mathbf{X}\|_{2p}^{2p} < \infty$ is strong enough to justify using the pull-through property in this step.

To continue, we must find a copy of E within the latter expression. We can accomplish this goal using one of the basic results from the theory of Schatten norms [Bha97, Cor. IV.2.6].

Proposition 7.6 (Hölder Inequality for Trace). *Let p and q be Hölder conjugate indices, i.e., positive numbers with the relationship $q = p/(p-1)$. Then*

$$\mathrm{tr}(\mathbf{BC}) \leq \|\mathbf{B}\|_p \|\mathbf{C}\|_q \quad \text{for all } \mathbf{B}, \mathbf{C} \in \mathbb{M}^d.$$

To complete the argument, apply the Hölder inequality for the trace followed by the Hölder inequality for the expectation. Thus,

$$\begin{aligned} E &\leq (2p-1) \cdot \mathbb{E} \left[\|\Delta_{\mathbf{X}}\|_p \cdot \|\mathbf{X}\|_{p/(p-1)}^{2p-2} \right] \\ &= (2p-1) \cdot \mathbb{E} \left[\|\Delta_{\mathbf{X}}\|_p \cdot \|\mathbf{X}\|_{2p}^{2p-2} \right] \\ &\leq (2p-1) \cdot (\mathbb{E} \|\Delta_{\mathbf{X}}\|_p^p)^{1/p} \cdot (\mathbb{E} \|\mathbf{X}\|_{2p}^{2p})^{(p-1)/p} \\ &= (2p-1) \cdot (\mathbb{E} \|\Delta_{\mathbf{X}}\|_p^p)^{1/p} \cdot E^{(p-1)/p}. \end{aligned}$$

Solve this algebraic inequality for the positive number E to conclude that

$$E \leq (2p-1)^p \cdot \mathbb{E} \|\Delta_{\mathbf{X}}\|_p^p.$$

Extract the $(2p)$ th root to establish the matrix BDG inequality.

8. EXTENSION TO GENERAL COMPLEX MATRICES

Although it may seem that our theory is limited to random Hermitian matrices, results for general random matrices follow as a formal corollary [Rec11, Tro11b]. The approach is based on a device from operator theory [Pau02].

Definition 8.1 (Hermitian Dilation). Let \mathbf{B} be a matrix in $\mathbb{C}^{d_1 \times d_2}$, and set $d = d_1 + d_2$. The *Hermitian dilation* of \mathbf{B} is the matrix

$$\mathcal{D}(\mathbf{B}) := \begin{bmatrix} \mathbf{0} & \mathbf{B} \\ \mathbf{B}^* & \mathbf{0} \end{bmatrix} \in \mathbb{H}^d. \quad (8.1)$$

The dilation has two valuable properties. First, it preserves spectral information:

$$\lambda_{\max}(\mathcal{D}(\mathbf{B})) = \|\mathcal{D}(\mathbf{B})\| = \|\mathbf{B}\|. \quad (8.2)$$

Second, the square of the dilation satisfies

$$\mathcal{D}(\mathbf{B})^2 = \begin{bmatrix} \mathbf{B}\mathbf{B}^* & \mathbf{0} \\ \mathbf{0} & \mathbf{B}^*\mathbf{B} \end{bmatrix}. \quad (8.3)$$

We can study a random matrix—not necessarily Hermitian—by applying our matrix concentration inequalities to the Hermitian dilation of the random matrix. As an illustration, let us prove a Bernstein inequality for general random matrices.

Corollary 8.2 (Bernstein Inequality for General Matrices). *Consider a finite sequence $(\mathbf{Z}_k)_{k \geq 1}$ of independent random matrices in $\mathbb{C}^{d_1 \times d_2}$ that satisfy*

$$\mathbb{E} \mathbf{Z}_k = \mathbf{0} \quad \text{and} \quad \|\mathbf{Z}_k\| \leq R \quad \text{for each index } k.$$

Define $d := d_1 + d_2$, and introduce the variance measure

$$\sigma^2 := \max \left\{ \left\| \sum_k \mathbb{E}(\mathbf{Z}_k \mathbf{Z}_k^*) \right\|, \left\| \sum_k \mathbb{E}(\mathbf{Z}_k^* \mathbf{Z}_k) \right\| \right\}.$$

Then, for all $t \geq 0$,

$$\mathbb{P} \left\{ \left\| \sum_k \mathbf{Z}_k \right\| \geq t \right\} \leq d \cdot \exp \left\{ \frac{-t^2}{3\sigma^2 + 2Rt} \right\}.$$

Furthermore,

$$\mathbb{E} \left\| \sum_k \mathbf{Z}_k \right\| \leq \sigma \sqrt{3 \log d} + R \log d.$$

Proof. Consider the random series $\sum_k \mathcal{D}(\mathbf{Z}_k)$. The summands are independent, random Hermitian matrices that satisfy

$$\mathbb{E} \mathcal{D}(\mathbf{Z}_k) = \mathbf{0} \quad \text{and} \quad \|\mathcal{D}(\mathbf{Z}_k)\| \leq R.$$

The second identity depends on the spectral property (8.2). Therefore, the matrix Bernstein inequality, Corollary 5.2, applies. To state the outcome, we first note that $\lambda_{\max}(\sum_k \mathcal{D}(\mathbf{Z}_k)) = \|\sum_k \mathbf{Z}_k\|$, again because of the spectral property (8.2). Next, use the formula (8.3) to compute that

$$\left\| \sum_k \mathbb{E} [\mathcal{D}(\mathbf{Z}_k)^2] \right\| = \left\| \begin{bmatrix} \sum_k \mathbb{E}(\mathbf{Z}_k \mathbf{Z}_k^*) & \mathbf{0} \\ \mathbf{0} & \sum_k \mathbb{E}(\mathbf{Z}_k^* \mathbf{Z}_k) \end{bmatrix} \right\| = \sigma^2.$$

This observation completes the proof. \square

9. THE SUM OF CONDITIONALLY ZERO-MEAN MATRICES

A chief advantage of the method of exchangeable pairs is its ability to handle random matrices constructed from *dependent* random variables. In this section, we briefly describe a way to relax the independence requirement when studying a sum of random matrices. In Sections 10 and 11, we develop more elaborate examples.

9.1. Formulation. Consider a finite sequence $(\mathbf{Y}_1, \dots, \mathbf{Y}_n)$ of random Hermitian matrices. We say that the sequence has the *conditional zero-mean property* when

$$\mathbb{E}[\mathbf{Y}_k \mid (\mathbf{Y}_j)_{j \neq k}] = \mathbf{0} \quad \text{almost surely for each index } k. \quad (9.1)$$

This definition is related to the conditional expectation property of a martingale difference sequence, although it is more restrictive. Suppose that we are interested in a sum of conditionally zero-mean random matrices:

$$\mathbf{X} := \mathbf{Y}_1 + \dots + \mathbf{Y}_n. \quad (9.2)$$

This type of series is quite common because it includes the case of a Rademacher series with random matrix coefficients.

Example 9.1 (Rademacher Series with Random Matrix Coefficients). Consider a finite sequence $(\mathbf{W}_k)_{k \geq 1}$ of random Hermitian matrices. Suppose that the sequence $(\varepsilon_k)_{k \geq 1}$ consists of independent Rademacher random variables that are independent from the random matrices. Consider the random series

$$\sum_k \varepsilon_k \mathbf{W}_k.$$

The summands may be strongly dependent on each other, but the independence of the Rademacher variables ensures that the summands satisfy the conditional zero-mean property (9.1).

9.2. A Matrix Stein Pair. Let us describe how to build a matrix Stein pair $(\mathbf{X}, \mathbf{X}')$ for the sum (9.2) of conditionally zero-mean random matrices. The approach is similar to the case of an independent sum, which appears in Section 2.4. For each k , we draw a random matrix \mathbf{Y}'_k so that \mathbf{Y}'_k and \mathbf{Y}_k are conditionally i.i.d. given $(\mathbf{Y}_j)_{j \neq k}$. Then, independently, we draw an index K uniformly at random from $\{1, \dots, n\}$. As in Section 2.4, the random matrix

$$\mathbf{X}' := \mathbf{Y}_1 + \dots + \mathbf{Y}_{K-1} + \mathbf{Y}'_K + \mathbf{Y}_{K+1} + \dots + \mathbf{Y}_n$$

is an exchangeable counterpart to \mathbf{X} . The conditional zero-mean property (9.1) ensures that

$$\begin{aligned} \mathbb{E}[\mathbf{X} - \mathbf{X}' \mid (\mathbf{Y}_j)_{j \geq 1}] &= \mathbb{E}[\mathbf{Y}_K - \mathbf{Y}'_K \mid (\mathbf{Y}_j)_{j \geq 1}] \\ &= \frac{1}{n} \sum_{k=1}^n (\mathbf{Y}_k - \mathbb{E}[\mathbf{Y}'_k \mid (\mathbf{Y}_j)_{j \neq k}]) = \frac{1}{n} \sum_{k=1}^n \mathbf{Y}_k = \frac{1}{n} \mathbf{X}. \end{aligned}$$

Therefore, $(\mathbf{X}, \mathbf{X}')$ is a matrix Stein pair with scale factor $\alpha = n^{-1}$.

We can determine the conditional variance after a short argument that parallels the computation (2.6) in the independent setting:

$$\Delta_{\mathbf{X}} = \frac{n}{2} \cdot \mathbb{E}[(\mathbf{Y}_K - \mathbf{Y}'_K)^2 \mid (\mathbf{Y}_j)_{j \geq 1}] = \frac{1}{2} \sum_{k=1}^n (\mathbf{Y}_k^2 + \mathbb{E}[\mathbf{Y}_k^2 \mid (\mathbf{Y}_j)_{j \neq k}]). \quad (9.3)$$

The expression (9.3) shows that, even in the presence of some dependence, we can control the size of the conditional expectation uniformly if we control the size of the individual summands.

Using the Stein pair $(\mathbf{X}, \mathbf{X}')$ and the expression (9.3), we may develop a variety of concentration inequalities for conditionally zero-mean sums that are analogous to our results for independent sums. We omit detailed examples.

10. COMBINATORIAL SUMS OF MATRICES

The method of exchangeable pairs can also be applied to many types of highly symmetric distributions. In this section, we study a class of *combinatorial matrix statistics*, which generalize the scalar statistics studied by Hoeffding [Hoe51].

10.1. Formulation. Consider a deterministic array $(\mathbf{A}_{jk})_{j,k=1}^n$ of Hermitian matrices, and let π be a uniformly random permutation on $\{1, \dots, n\}$. Define the random matrix

$$\mathbf{Y} := \sum_{j=1}^n \mathbf{A}_{j\pi(j)} \quad \text{whose mean} \quad \mathbb{E} \mathbf{Y} = \frac{1}{n} \sum_{j,k=1}^n \mathbf{A}_{jk}. \quad (10.1)$$

The combinatorial sum \mathbf{Y} is a natural candidate for an exchangeable pair analysis because the random permutation is highly symmetric. Before we describe how to construct a matrix Stein pair, let us mention a few problems that lead to a random matrix of the form \mathbf{Y} .

Example 10.1 (Sampling without Replacement). Consider a finite collection $\mathcal{B} := \{\mathbf{B}_1, \dots, \mathbf{B}_n\}$ of deterministic Hermitian matrices. Suppose that we want to study a sum of s matrices sampled randomly from \mathcal{B} without replacement. We can express this type of series in the form

$$\mathbf{W} := \sum_{j=1}^s \mathbf{B}_{\pi(j)},$$

where π is a random permutation on $\{1, \dots, n\}$. The matrix \mathbf{W} is therefore an example of a combinatorial sum.

Example 10.2 (A Randomized “Inner Product”). Consider two fixed sequences of complex matrices

$$\mathbf{B}_1, \dots, \mathbf{B}_n \in \mathbb{C}^{d_1 \times s} \quad \text{and} \quad \mathbf{C}_1, \dots, \mathbf{C}_n \in \mathbb{C}^{s \times d_2}.$$

We may form a permuted matrix “inner product” by arranging one sequence in random order, multiplying the elements of the two sequences together, and summing the terms. In other words, we are interested in the random matrix

$$\mathbf{Z} := \sum_{j=1}^n \mathbf{B}_j \mathbf{C}_{\pi(j)}.$$

Then the random matrix $\mathcal{D}(\mathbf{Z})$ is a combinatorial sum of Hermitian matrices.

10.2. A Matrix Stein Pair. To study the combinatorial sum (10.1) of matrices using the method of exchangeable pairs, we first introduce the zero-mean random matrix

$$\mathbf{X} := \mathbf{Y} - \mathbb{E} \mathbf{Y}.$$

To construct a matrix Stein pair $(\mathbf{X}, \mathbf{X}')$, we draw a pair (J, K) of indices independently of π and uniformly at random from $\{1, \dots, n\}^2$. Define a second random permutation $\pi' := \pi \circ (J, K)$ by composing π with the transposition of the random indices J and K . The pair (π, π') is exchangeable, so the matrix

$$\mathbf{X}' := \sum_{j=1}^n \mathbf{A}_{j\pi'(j)} - \mathbb{E} \mathbf{Y}$$

is an exchangeable counterpart to \mathbf{X} .

To verify that $(\mathbf{X}, \mathbf{X}')$ is a matrix Stein pair, we calculate that

$$\begin{aligned} \mathbb{E}[\mathbf{X} - \mathbf{X}' \mid \pi] &= \mathbb{E}[\mathbf{A}_{J\pi(J)} + \mathbf{A}_{K\pi(K)} - \mathbf{A}_{J\pi(K)} - \mathbf{A}_{K\pi(J)} \mid \pi] \\ &= \frac{1}{n^2} \sum_{j,k=1}^n [\mathbf{A}_{j\pi(j)} + \mathbf{A}_{k\pi(k)} - \mathbf{A}_{j\pi(k)} - \mathbf{A}_{k\pi(j)}] = \frac{2}{n}(\mathbf{Y} - \mathbb{E} \mathbf{Y}) = \frac{2}{n} \mathbf{X}. \end{aligned}$$

The first identity holds because the random sums \mathbf{X} and \mathbf{X}' differ for only four choices of indices. We see that $(\mathbf{X}, \mathbf{X}')$ is a Stein pair with scale factor $\alpha = 2/n$.

Turning to the conditional variance, we find that

$$\Delta_{\mathbf{X}}(\pi) = \frac{n}{4} \mathbb{E}[(\mathbf{X} - \mathbf{X}')^2 \mid \pi] = \frac{1}{4n} \sum_{j,k=1}^n [\mathbf{A}_{j\pi(j)} + \mathbf{A}_{k\pi(k)} - \mathbf{A}_{j\pi(k)} - \mathbf{A}_{k\pi(j)}]^2. \quad (10.2)$$

The structure of the conditional variance is somewhat different from previous examples, but we recognize that $\Delta_{\mathbf{X}}$ is controlled when the matrices \mathbf{A}_{jk} are bounded.

10.3. Exponential Concentration for a Combinatorial Sum. We can apply our matrix concentration results to study the behavior of a combinatorial sum of matrices. As an example, let us present a Bernstein-type inequality. The argument is similar to the proof of Corollary 5.2, so we leave the details to Appendix C.

Corollary 10.3 (Bernstein Inequality for a Combinatorial Sum of Matrices). *Consider an array $(\mathbf{A}_{jk})_{j,k=1}^n$ of deterministic matrices in \mathbb{H}^d that satisfy*

$$\sum_{j,k=1}^n \mathbf{A}_{jk} = \mathbf{0} \quad \text{and} \quad \|\mathbf{A}_{jk}\| \leq R \quad \text{for each pair } (j, k) \text{ of indices.}$$

Define the random matrix $\mathbf{X} := \sum_{j=1}^n \mathbf{A}_{j\pi(j)}$, where π is a uniformly random permutation on $\{1, \dots, n\}$. Then, for all $t \geq 0$,

$$\mathbb{P}\{\lambda_{\max}(\mathbf{X}) \geq t\} \leq d \cdot \exp\left\{\frac{-t^2}{12\sigma^2 + 4\sqrt{2}Rt}\right\} \quad \text{where} \quad \sigma^2 := \frac{1}{n} \left\| \sum_{j,k=1}^n \mathbf{A}_{jk}^2 \right\|.$$

Furthermore,

$$\mathbb{E} \lambda_{\max}(\mathbf{X}) \leq \sigma \sqrt{12 \log d} + 2\sqrt{2} R \log d.$$

11. SELF-REPRODUCING MATRIX FUNCTIONS

The method of exchangeable pairs can also be used to analyze nonlinear matrix-valued functions of random variables. In this section, we explain how to analyze matrix functions that satisfy a *self-reproducing property*.

11.1. Example: Matrix Second-Order Rademacher Chaos. We begin with an example that shows how the self-reproducing property might arise. Consider a quadratic form that takes on random matrix values:

$$\mathbf{H}(\varepsilon) := \sum_k \sum_{j < k} \varepsilon_j \varepsilon_k \mathbf{A}_{jk}. \quad (11.1)$$

In this expression, ε is a finite vector of independent Rademacher random variables. The array $(\mathbf{A}_{jk})_{j,k \geq 1}$ consists of deterministic Hermitian matrices, and we assume that $\mathbf{A}_{jk} = \mathbf{A}_{kj}$.

Observe that the summands in $\mathbf{H}(\varepsilon)$ are dependent and that they need not satisfy the conditional zero-mean property (9.1). Nevertheless, $\mathbf{H}(\varepsilon)$ does satisfy a fruitful self-reproducing property:

$$\begin{aligned} \sum_k (\mathbf{H}(\varepsilon) - \mathbb{E}[\mathbf{H}(\varepsilon) | (\varepsilon_j)_{j \neq k}]) &= \sum_k \sum_{j \neq k} \varepsilon_j (\varepsilon_k - \mathbb{E}[\varepsilon_k]) \mathbf{A}_{jk} \\ &= \sum_k \sum_{j \neq k} \varepsilon_j \varepsilon_k \mathbf{A}_{jk} = 2\mathbf{H}(\varepsilon). \end{aligned}$$

We have applied the pull-through property of conditional expectation, the assumption that the Rademacher variables are independent, and the fact that $\mathbf{A}_{jk} = \mathbf{A}_{kj}$. As we will see, this type of self-reproducing property can be used to construct a matrix Stein pair.

A random matrix of the form (11.1) is called a *second-order Rademacher chaos*. This class of random matrices arises in a variety of situations, including randomized linear algebra [CD11a], compressed sensing [Rau10, Sec. 9], and chance-constrained optimization [CSW11]. Indeed, concentration inequalities for the matrix-valued Rademacher chaos have a wide range of potential applications.

11.2. Formulation and Matrix Stein Pair. In this section, we describe a more general version of the self-reproducing property. Suppose that $\mathbf{z} := (Z_1, \dots, Z_n)$ is a random vector taking values in a Polish space \mathcal{Z} . First, we construct an exchangeable counterpart

$$\mathbf{z}' := (Z_1, \dots, Z_{K-1}, Z'_K, Z_{K+1}, \dots, Z_n) \quad (11.2)$$

where Z_k and Z'_k are conditionally i.i.d. given $(Z_j)_{j \neq k}$ and K is an independent coordinate drawn uniformly at random from $\{1, \dots, n\}$.

Next, let $\mathbf{H} : \mathcal{Z}^n \rightarrow \mathbb{H}^d$ be a bounded measurable function. Assume that $\mathbf{H}(\mathbf{z})$ satisfies an abstract *self-reproducing property*: for a parameter $s > 0$,

$$\sum_{k=1}^n (\mathbf{H}(\mathbf{z}) - \mathbb{E}[\mathbf{H}(Z_1, \dots, Z'_k, \dots, Z_n) | \mathbf{z}]) = s \cdot (\mathbf{H}(\mathbf{z}) - \mathbb{E} \mathbf{H}(\mathbf{z})) \quad \text{almost surely.} \quad (11.3)$$

Under this assumption, we can easily check that the random matrices

$$\mathbf{X} := \mathbf{H}(\mathbf{z}) - \mathbb{E} \mathbf{H}(\mathbf{z}) \quad \text{and} \quad \mathbf{X}' := \mathbf{H}(\mathbf{z}') - \mathbb{E} \mathbf{H}(\mathbf{z})$$

form a matrix Stein pair. Indeed,

$$\mathbb{E}[\mathbf{X} - \mathbf{X}' | \mathbf{z}] = \mathbb{E}[\mathbf{H}(\mathbf{z}) - \mathbf{H}(\mathbf{z}') | \mathbf{z}] = \frac{s}{n} (\mathbf{H}(\mathbf{z}) - \mathbb{E} \mathbf{H}(\mathbf{z})) = \frac{s}{n} \mathbf{X}.$$

We determine that $(\mathbf{X}, \mathbf{X}')$ is a matrix Stein pair with scaling factor $\alpha = s/n$.

Finally, we compute the conditional variance:

$$\Delta_{\mathbf{X}}(\mathbf{z}) = \frac{n}{2s} \mathbb{E}[(\mathbf{H}(\mathbf{z}) - \mathbf{H}(\mathbf{z}'))^2 | \mathbf{z}] = \frac{1}{2s} \sum_{k=1}^n \mathbb{E}[(\mathbf{H}(\mathbf{z}) - \mathbf{H}(Z_1, \dots, Z'_k, \dots, Z_n))^2 | \mathbf{z}].$$

We discover that the conditional variance is small when each coordinate of \mathbf{H} has controlled variance. In this case, the method of exchangeable pairs provides good concentration inequalities for the random matrix \mathbf{X} .

11.3. Matrix Bounded Differences Inequality. As an example of this framework, we can develop a bounded differences inequality for random matrices by appealing to Theorem 4.1.

Corollary 11.1 (Matrix Bounded Differences). *Let $\mathbf{z} := (Z_1, \dots, Z_n)$ be a random vector taking values in a Polish space \mathcal{Z} , and let \mathbf{z}' be an exchangeable counterpart, constructed as in (11.2). Suppose that $\mathbf{H} : \mathcal{Z}^n \rightarrow \mathbb{H}^d$ is a bounded function that satisfies the self-reproducing property*

$$\sum_{k=1}^n (\mathbf{H}(\mathbf{z}) - \mathbb{E}[\mathbf{H}(Z_1, \dots, Z'_k, \dots, Z_n) | \mathbf{z}]) = s \cdot (\mathbf{H}(\mathbf{z}) - \mathbb{E} \mathbf{H}(\mathbf{z})) \quad \text{almost surely}$$

for a parameter $s > 0$ as well as the bounded differences condition

$$\mathbb{E} [(\mathbf{H}(\mathbf{z}) - \mathbf{H}(Z_1, \dots, Z'_k, \dots, Z_n))^2 | \mathbf{z}] \preceq \mathbf{A}_k^2 \quad \text{almost surely for each index } k, \quad (11.4)$$

where \mathbf{A}_k is a deterministic matrix in \mathbb{H}^d . Then, for all $t \geq 0$,

$$\mathbb{P} \{ \lambda_{\max}(\mathbf{H}(\mathbf{z}) - \mathbb{E} \mathbf{H}(\mathbf{z})) \geq t \} \leq d \cdot e^{-st^2/L} \quad \text{where } L := \left\| \sum_{k=1}^n \mathbf{A}_k^2 \right\|.$$

Furthermore,

$$\mathbb{E} \lambda_{\max}(\mathbf{H}(\mathbf{z}) - \mathbb{E} \mathbf{H}(\mathbf{z})) \leq \sqrt{\frac{L \log d}{s}}.$$

In the scalar setting, Corollary 11.1 reduces to a version of McDiarmid's bounded difference inequality [McD89]. The result also complements the matrix bounded difference inequality of [Tro11b, Cor. 7.5], which requires independent input variables but makes no self-reproducing assumption.

Proof. Since $\mathbf{H}(\mathbf{z})$ is self-reproducing, we may construct a matrix Stein pair $(\mathbf{X}, \mathbf{X}')$ with scale factor $\alpha = s/n$ as in Section 11. According to (11.2), the conditional variance of the pair satisfies

$$\begin{aligned} \Delta_{\mathbf{X}} &= \frac{1}{2s} \sum_{k=1}^n \mathbb{E} [(\mathbf{H}(\mathbf{z}) - \mathbf{H}(Z_1, \dots, Z'_k, \dots, Z_n))^2 | \mathbf{z}] \\ &\preceq \frac{1}{2s} \sum_{k=1}^n \mathbf{A}_k^2 \preceq \frac{L}{2s} \cdot \mathbf{I}. \end{aligned}$$

We have used the bounded differences condition (11.4) and the definition of the bound L . To complete the proof, we apply the concentration result, Theorem 4.1, with the parameters $c = 0$ and $v = L/2s$. \square

12. EXTENSIONS AND FUTURE WORK

Although the examples we treat in this work all involve matrix Stein pairs, we are aware of problems that require a larger class of exchangeable pairs. The following definition represents a significant extension of the matrix Stein pair formalism.

Definition 12.1 (Generalized Matrix Stein Pair). Let $g : \mathbb{R} \rightarrow \mathbb{R}$ be a weakly increasing function. Let (Z, Z') be an exchangeable pair of random variables taking values in a Polish space \mathcal{Z} , and let $\Psi : \mathcal{Z} \rightarrow \mathbb{H}^d$ be a measurable function. Define the random Hermitian matrices

$$\mathbf{X} := \Psi(Z) \quad \text{and} \quad \mathbf{X}' := \Psi(Z').$$

We say that $(\mathbf{X}, \mathbf{X}')$ is a *generalized matrix Stein pair* if there is a constant $\alpha \in (0, 1]$ for which

$$\mathbb{E}[g(\mathbf{X}) - g(\mathbf{X}') | Z] = \alpha \mathbf{X} \quad \text{almost surely.}$$

When the function g is linear, a generalized matrix Stein pair reduces to a matrix Stein pair. The *conditional variance* of the generalized pair is defined as

$$\Delta_{\mathbf{X}} := \Delta_{\mathbf{X}}(Z) := \frac{1}{2\alpha} \cdot \text{Re} \mathbb{E} [(g(\mathbf{X}) - g(\mathbf{X}')) \cdot (\mathbf{X} - \mathbf{X}') | Z],$$

where $\text{Re}(\mathbf{A}) := \frac{1}{2}(\mathbf{A} + \mathbf{A}^*)$ refers to the Hermitian part of a square matrix.

We have chosen to focus on the simpler definition of a Stein pair because the presentation is more transparent. Nevertheless, the results in this paper all extend readily to generalized matrix Stein pairs with little extra work. We leave applications of generalized matrix Stein pairs for future research.

ACKNOWLEDGMENTS

Some of this material is based upon work supported in part by the U. S. Army Research Laboratory and the U. S. Army Research Office under contract/grant number W911NF-11-1-0391. Mackey was supported by the National Defense Science and Engineering Graduate Fellowship. Chen, Farrell, and Tropp were supported by ONR awards N00014-08-1-0883 and N00014-11-1002, AFOSR award FA9550-09-1-0643, DARPA award N66001-08-1-2065, and a Sloan Research Fellowship. The authors thank Houman Owahdi for helpful conversations.

REFERENCES

- [AW02] R. Ahlswede and A. Winter. Strong converse for identification via quantum channels. *IEEE Trans. Inform. Theory*, 48(3):569–579, Mar. 2002.
- [Bha97] R. Bhatia. *Matrix Analysis*. Number 169 in Graduate Texts in Mathematics. Springer, Berlin, 1997.
- [Buc01] A. Buchholz. Operator Khintchine inequality in non-commutative probability. *Math. Ann.*, 319:1–16, 2001.
- [Bur73] D. L. Burkholder. Distribution function inequalities for martingales. *Ann. Probab.*, 1:19–42, 1973.
- [Car10] E. Carlen. Trace inequalities and quantum entropy: an introductory course. In *Entropy and the quantum: Arizona School of Analysis with Applications, March 16-20, 2009, University of Arizona*, Contemporary Mathematics. American Mathematical Society, Providence, RI, 2010.
- [CD11a] J. Chiu and L. Demanet. Matrix probing and its conditioning. Available at [arXiv:1102.5689](#), Feb. 2011.
- [CD11b] J. Chiu and L. Demanet. Sublinear randomized algorithms for skeleton decomposition. Available at [arXiv:1110.4193](#), Oct. 2011.
- [CDL11] A. Cohen, M. Davenport, and D. Leviatan. On the stability and accuracy of least-squares approximation. Available at [arXiv:1111.4422](#), Nov. 2011.
- [CGT12] R. Y. Chen, A. Gittens, and J. A. Tropp. The masked sample covariance estimator: An analysis via the matrix Laplace transform method. *Information and Inference*, 2012. To appear.
- [Cha07] S. Chatterjee. Stein’s method for concentration inequalities. *Probab. Theory Related Fields*, 138:305–321, 2007.
- [Cha08] S. Chatterjee. *Concentration inequalities with exchangeable pairs*. PhD thesis, Stanford University, Palo Alto, Feb. 2008.
- [CSW11] S.-S. Cheung, A. Man-Cho So, and K. Wang. Chance-constrained linear matrix inequalities with dependent perturbations: A safe tractable approximation approach. Available at http://www.se.cuhk.edu.hk/~manchoso/papers/cclmi_sta.pdf, 2011.
- [Git11] A. Gittens. The spectral norm error of the naïve Nyström extension. Available at [arXiv:1110.5305](#), Oct. 2011.
- [Gro11] D. Gross. Recovering low-rank matrices from few coefficients in any basis. *IEEE Trans. Inform. Theory*, 57(3):1548–1566, Mar. 2011.
- [GT11] A. Gittens and J. A. Tropp. Tail bounds for all eigenvalues of a sum of random matrices. Available at [arXiv:1104.4513](#), April 2011.
- [Hig08] N. J. Higham. *Functions of Matrices: Theory and Computation*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2008.
- [HKZ11] D. Hsu, S. M. Kakade, and T. Zhang. Dimension-free tail inequalities for sums of random matrices. Available at [arXiv:1104.1672](#), 2011.
- [Hoe51] W. Hoeffding. A combinatorial central limit theorem. *Ann. Math. Statist.*, 22:558–566, 1951.
- [Hoe63] W. Hoeffding. Probability inequalities for sums of bounded random variables. *J. Amer. Statist. Assoc.*, 58(301):13–30, 1963.
- [JX03] M. Junge and Q. Xu. Noncommutative Burkholder/Rosenthal inequalities. *Ann. Probab.*, 31(2):948–995, 2003.
- [JX08] M. Junge and Q. Xu. Noncommutative Burkholder/Rosenthal inequalities II: Applications. *Israel J. Math.*, 167:227–282, 2008.
- [JZ11] M. Junge and Q. Zheng. Noncommutative Bennett and Rosenthal inequalities. Available at [arXiv:1111.1027](#), Nov. 2011.

- [Kol11] V. Koltchinskii. *Oracle Inequalities in Empirical Risk Minimization and Sparse Recovery Problems*. Number XXXVIII-2008 in Ecole d'Eté de Probabilités de Saint-Flour. Springer, 2011.
- [Led01] M. Ledoux. *The Concentration of Measure Phenomenon*. American Mathematical Society, Providence, RI, 2001.
- [Lie73] E. H. Lieb. Convex trace functions and the Wigner–Yanase–Dyson conjecture. *Adv. Math.*, 11:267–288, 1973.
- [LP86] F. Lust-Piquard. Inégalités de Khintchine dans C_p ($1 < p < \infty$). *C. R. Math. Acad. Sci. Paris*, 303(7):289–292, 1986.
- [LPP91] F. Lust-Piquard and G. Pisier. Noncommutative Khintchine and Paley inequalities. *Ark. Mat.*, 29(2):241–260, 1991.
- [McD89] C. McDiarmid. On the method of bounded differences. In *Surveys in Combinatorics 1989*, pages 148–188, London, 1989. London Mathematical Society Lecture Notes.
- [MTJ11] L. Mackey, A. Talwalkar, and M. I. Jordan. Divide-and-conquer matrix factorization. In *Advances in Neural Information Processing Systems 24*. MIT Press, Granada, Dec. 2011.
- [NP77] S. V. Nagaev and I. F. Pinelis. Some inequalities for the distribution of sums of independent random variables. *Theory Probab. Appl.*, 22(2):248–256, 1977.
- [Oli09] R. I. Oliveira. Concentration of the adjacency matrix and of the Laplacian in random graphs with independent edges. Available at [arXiv:0911.0600](https://arxiv.org/abs/0911.0600), Nov. 2009.
- [Oli10] R. I. Oliveira. Sums of random Hermitian matrices and an inequality by Rudelson. *Electron. Commun. Probab.*, 15:203–212, 2010.
- [Pau02] V. I. Paulsen. *Completely Bounded Maps and Operator Algebras*. Number 78 in Cambridge Studies in Advanced Mathematics. Cambridge Univ. Press, Cambridge, 2002.
- [Pet94] D. Petz. A survey of certain trace inequalities. In *Functional analysis and operator theory*, volume 30 of *Banach Center Publications*, pages 287–298, Warsaw, 1994. Polish Acad. Sci.
- [Pin94] I. F. Pinelis. Optimum bounds for the distributions of martingales in Banach spaces. *Ann. Probab.*, 22(4):1679–1706, 1994.
- [PX97] G. Pisier and Q. Xu. Non-commutative martingale inequalities. *Comm. Math. Phys.*, 189(3):667–698, 1997.
- [Rau10] H. Rauhut. Compressive sensing and structured random matrices. In *Theoretical Foundations and Numerical Methods for Sparse Recovery*, volume 9 of *Radon Ser. Comp. Appl. Math.*, pages 1–93. de Gruyter, Berlin, 2010.
- [Rec11] B. Recht. Simpler approach to matrix completion. *J. Mach. Learn. Res.*, 12:3413–3430, 2011.
- [Ros70] H. P. Rosenthal. On subspaces of L_p ($p > 2$) spanned by sequences of independent random variables. *Israel J. Math.*, 8:273–303, 1970.
- [Rud99] M. Rudelson. Random vectors in the isotropic position. *J. Funct. Anal.*, 164:60–72, 1999.
- [RV07] M. Rudelson and R. Vershynin. Sampling from large matrices: An approach through geometric functional analysis. *J. Assoc. Comput. Mach.*, 54(4):Article 21, 19 pp., Jul. 2007. (electronic).
- [So11] A. Man-Cho So. Moment inequalities for sums of random matrices and their applications in optimization. *Math. Program.*, 130(1):125–151, 2011.
- [Ste72] C. Stein. A bound for the error in the normal approximation to the distribution of a sum of dependent random variables. In *Proc. 6th Berkeley Symp. Math. Statist. Probab.*, Berkeley, 1972. Univ. California Press.
- [Tro11a] J. A. Tropp. Freedman’s inequality for matrix martingales. *Electron. Commun. Probab.*, 16:262–270, 2011.
- [Tro11b] J. A. Tropp. User-friendly tail bounds for sums of random matrices. *Found. Comput. Math.*, August 2011.

APPENDIX A. PROOF OF THEOREM 5.1: REFINED EXPONENTIAL CONCENTRATION

The proof of Theorem 5.1 parallels the argument in Theorem 4.1, but it differs at an important point. In the earlier result, we used an almost-sure bound on the conditional variance to control the derivative of the trace mgf. This time, we use entropy inequalities to introduce finer information about the behavior of the conditional variance. The proof is essentially a matrix version of Chatterjee’s argument [Cha08, Thm. 3.13].

Let $(\mathbf{X}, \mathbf{X}')$ be a matrix Stein pair. Consider the normalized trace mgf

$$m(\theta) := \mathbb{E} \bar{\text{tr}} e^{\theta \mathbf{X}} \quad \text{for } \theta \geq 0. \quad (\text{A.1})$$

Our main object is to establish the following inequality for the trace mgf.

$$\log m(\theta) \leq \frac{r(\psi) \theta^2}{2(1 - \theta^2/\psi)} \quad \text{for } 0 \leq \theta < \sqrt{\psi} \text{ and } \psi > 0. \quad (\text{A.2})$$

The function $r(\psi)$ is defined in (5.1). We establish (A.2) in Section A.1 et seq. Afterward, in Section A.5, we invoke the matrix Laplace transform bound to complete the proof of Theorem 5.1.

A.1. The Derivative of the Trace Mgf. The first steps of the argument are the same as in the proof of Theorem 4.1. Since \mathbf{X} is almost surely bounded, we need not worry about regularity conditions. The derivative of the trace mgf satisfies

$$m'(\theta) = \mathbb{E} \operatorname{tr} [\mathbf{X} e^{\theta \mathbf{X}}] \quad \text{for } \theta \in \mathbb{R}. \quad (\text{A.3})$$

Lemma 3.7 provides a bound for the derivative in terms of the conditional variance:

$$m'(\theta) \leq \theta \cdot \mathbb{E} \bar{\operatorname{tr}} [\Delta_{\mathbf{X}} e^{\theta \mathbf{X}}] \quad \text{for } \theta \geq 0. \quad (\text{A.4})$$

In the proof of Lemma 4.3, we applied an almost-sure bound for the conditional variance to control the derivative of the mgf. This time, we incorporate information about the typical size of $\Delta_{\mathbf{X}}$ by developing a bound in terms of the function $r(\psi)$.

A.2. Entropy for Random Matrices and Duality. Let us introduce an entropy function for random matrices.

Definition A.1 (Entropy for Random Matrices). Let \mathbf{W} be a random matrix in \mathbb{H}_+^d subject to the normalization $\mathbb{E} \bar{\operatorname{tr}} \mathbf{W} = 1$. The (negative) *matrix entropy* is defined as

$$\operatorname{ent}(\mathbf{W}) := \mathbb{E} \bar{\operatorname{tr}}(\mathbf{W} \log \mathbf{W}). \quad (\text{A.5})$$

We enforce the convention that $0 \log 0 = 0$.

The matrix entropy is relevant to our discussion because its Fenchel–Legendre conjugate is the cumulant generating function. The Young inequality for matrix entropy provides one way to formulate this duality relationship.

Proposition A.2 (Young Inequality for Matrix Entropy). *Suppose that \mathbf{V} is a random matrix in \mathbb{H}^d that is almost surely bounded in norm, and suppose that \mathbf{W} is a random matrix in \mathbb{H}_+^d subject to the normalization $\mathbb{E} \bar{\operatorname{tr}} \mathbf{W} = 1$. Then*

$$\mathbb{E} \bar{\operatorname{tr}}(\mathbf{V} \mathbf{W}) \leq \log \mathbb{E} \bar{\operatorname{tr}} e^{\mathbf{V}} + \operatorname{ent}(\mathbf{W}).$$

Proposition A.2 follows from an easy variant of the argument in [Car10, Thm. 2.13].

A.3. A Refined Differential Inequality for the Trace Mgf. We intend to apply the Young inequality for matrix entropy to decouple the product of random matrices in (A.4). First, we must rescale the exponential in (A.4) so its expected trace equals one:

$$\mathbf{W}(\theta) := \frac{1}{\mathbb{E} \bar{\operatorname{tr}} e^{\theta \mathbf{X}}} \cdot e^{\theta \mathbf{X}} = \frac{1}{m(\theta)} \cdot e^{\theta \mathbf{X}}. \quad (\text{A.6})$$

For each $\psi > 0$, we can rewrite (A.4) as

$$m'(\theta) \leq \frac{\theta m(\theta)}{\psi} \cdot \mathbb{E} \bar{\operatorname{tr}} [\psi \Delta_{\mathbf{X}} \cdot \mathbf{W}(\theta)].$$

The Young inequality for matrix entropy, Proposition A.2, implies that

$$m'(\theta) \leq \frac{\theta m(\theta)}{\psi} \left[\log \mathbb{E} \bar{\operatorname{tr}} e^{\psi \Delta_{\mathbf{X}}} + \operatorname{ent}(\mathbf{W}(\theta)) \right]. \quad (\text{A.7})$$

The first term in the bracket is precisely $\psi r(\psi)$. Let us examine the second term more closely.

To control the matrix entropy of $\mathbf{W}(\theta)$, we need to bound its logarithm. Referring back to the definition (A.6), we see that

$$\log \mathbf{W}(\theta) = \theta \mathbf{X} - (\log \mathbb{E} \bar{\operatorname{tr}} e^{\theta \mathbf{X}}) \cdot \mathbf{I} \preceq \theta \mathbf{X} - (\log \bar{\operatorname{tr}} e^{\theta \mathbb{E} \mathbf{X}}) \cdot \mathbf{I} = \theta \mathbf{X}. \quad (\text{A.8})$$

The second relation depends on Jensen's inequality and the fact that the trace exponential is convex [Pet94, Sec. 2]. The third relation relies on the property that $\mathbb{E} \mathbf{X} = \mathbf{0}$. Since the matrix $\mathbf{W}(\theta)$ is positive, we can substitute the semidefinite bound (A.8) into the definition (A.5) of the matrix entropy:

$$\text{ent}(\mathbf{W}(\theta)) = \mathbb{E} \bar{\text{tr}} [\mathbf{W}(\theta) \cdot \log \mathbf{W}(\theta)] \leq \theta \cdot \mathbb{E} \bar{\text{tr}} [\mathbf{W}(\theta) \cdot \mathbf{X}] = \frac{\theta}{m(\theta)} \cdot \mathbb{E} \bar{\text{tr}} [\mathbf{X} e^{\theta \mathbf{X}}].$$

We have reintroduced the definition (A.6) of $\mathbf{W}(\theta)$ in the last relation. Identify the derivative (A.3) of the trace mgf to reach

$$\text{ent}(\mathbf{W}(\theta)) \leq \frac{\theta m'(\theta)}{m(\theta)}. \quad (\text{A.9})$$

To establish a differential inequality, substitute the definition (5.1) of $r(\psi)$ and the bound (A.9) into the estimate (A.7) to discover that

$$m'(\theta) \leq \frac{\theta m(\theta)}{\psi} \left[\psi r(\psi) + \frac{\theta m'(\theta)}{m(\theta)} \right] = r(\psi) \theta \cdot m(\theta) + \frac{\theta^2}{\psi} \cdot m'(\theta).$$

Rearrange this formula to isolate the log-derivative $m'(\theta)/m(\theta)$ of the trace mgf. We conclude that

$$\frac{d}{d\theta} \log m(\theta) \leq \frac{r(\psi) \theta}{1 - \theta^2/\psi} \quad \text{for } 0 \leq \theta < \sqrt{\psi}. \quad (\text{A.10})$$

A.4. Solving the Differential Inequality. To integrate (A.10), recall that $\log m(0) = 0$, and invoke the fundamental theorem of calculus to see that

$$\log m(\theta) \leq \int_0^\theta \frac{r(\psi) s}{1 - s^2/\psi} ds \leq \int_0^\theta \frac{r(\psi) s}{1 - \theta^2/\psi} ds = \frac{r(\psi) \theta^2}{2(1 - \theta^2/\psi)}.$$

This calculation is valid whenever $0 \leq \theta < \sqrt{\psi}$. This is the claim (A.2).

A.5. The Matrix Laplace Transform Argument. With the trace mgf bound (A.2) at hand, we can complete the proof of Theorem 5.1. Proposition 3.3, the matrix Laplace transform method, yields the estimate

$$\mathbb{P} \{ \lambda_{\max}(\mathbf{X}) \geq t \} \leq d \cdot \inf_{0 < \theta < \sqrt{\psi}} \exp \left\{ -\theta t + \frac{\theta}{2} \cdot \frac{r(\psi) \theta}{1 - \theta^2/\psi} \right\} \leq d \cdot e^{-t \theta_\star(t)/2}, \quad (\text{A.11})$$

where we define $\theta_\star(t)$ implicitly as the positive root of the quadratic equation

$$\frac{r(\psi) \theta}{1 - \theta^2/\psi} = t.$$

Solve the quadratic equation to obtain the explicit formula

$$\theta_\star(t) = \frac{\psi r(\psi)}{2t} \left(\sqrt{1 + \frac{4t^2}{\psi r(\psi)^2}} - 1 \right). \quad (\text{A.12})$$

The numerical fact $\sqrt{1+a} < 1 + \sqrt{a}$, valid for $a > 0$, allows us to verify that $\theta_\star(t) < \sqrt{\psi}$. We can obtain a lower bound

$$\theta_\star(t) \geq \frac{t}{r(\psi) + t/\sqrt{\psi}} \quad (\text{A.13})$$

from the inequality $\sqrt{1+a} - 1 \geq a/(2 + \sqrt{a})$, which holds for $a \geq 0$. Indeed,

$$a = (\sqrt{1+a} - 1)(\sqrt{1+a} + 1) \leq (\sqrt{1+a} - 1)(2 + \sqrt{a}).$$

Introduce the estimate (A.13) into the probability inequality (A.11) to complete the proof of the tail bound (5.2).

To establish the inequality (5.3) for the expectation of the maximum eigenvalue, we can apply Proposition 3.3 and the trace mgf bound (A.2) a second time. Indeed,

$$\mathbb{E} \lambda_{\max}(\mathbf{X}) \leq \inf_{0 < \theta < \sqrt{\psi}} \left\{ \frac{\log d}{\theta} + \frac{1}{2} \cdot \frac{r(\psi) \theta}{1 - \theta^2 / \sqrt{\psi}} \right\} \leq \inf_{t \geq 0} \left\{ \frac{\log d}{\theta_*(t)} + \frac{t}{2} \right\},$$

where $\theta_*(t)$ is the function defined in (A.12). Incorporate the lower bound (A.13) to reach

$$\mathbb{E} \lambda_{\max}(\mathbf{X}) \leq \inf_{t \geq 0} \left\{ \frac{r(\psi) \log d}{t} + \frac{t}{2} + \frac{\log d}{\sqrt{\psi}} \right\} = \sqrt{2r(\psi) \log d} + \frac{\log d}{\sqrt{\psi}}.$$

This observation completes the proof of Theorem 5.1.

APPENDIX B. PROOF OF THEOREM 7.5: MATRIX ROSENTHAL INEQUALITY

The proof of the matrix Rosenthal inequality takes place in two steps. First, we verify that the bound (7.3) holds for psd random matrices. Then, we use this result to provide a short proof of the bound (7.4) for Hermitian random matrices. Before we start, let us remind the reader that the L_p norm of a scalar random variable Z is given by the expression $(\mathbb{E} |Z|^p)^{1/p}$ for each $p \geq 1$.

B.1. A Sum of Random Psd Matrices. We begin with the moment bound (7.3) for an independent sum of random psd matrices. Introduce the quantity of interest

$$E^2 := \left(\mathbb{E} \left\| \sum_k \mathbf{P}_k \right\|_{2p}^{2p} \right)^{1/2p}.$$

We may invoke the triangle inequality for the L_{2p} norm to obtain

$$E^2 \leq \left(\mathbb{E} \left\| \sum_k (\mathbf{P}_k - \mathbb{E} \mathbf{P}_k) \right\|_{2p}^{2p} \right)^{1/2p} + \left\| \sum_k \mathbb{E} \mathbf{P}_k \right\|_{2p} =: \left(\mathbb{E} \|\mathbf{X}\|_{2p}^{2p} \right)^{1/2p} + \mu.$$

We can apply the matrix BDG inequality to control this expectation, which yields an algebraic inequality between E^2 and E . We solve this inequality to bound E^2 .

The series \mathbf{X} consists of centered, independent random matrices, so we can use the Stein pair described in Section 2.4. According to (2.6), the conditional variance $\Delta_{\mathbf{X}}$ takes the form

$$\begin{aligned} \Delta_{\mathbf{X}} &= \frac{1}{2} \sum_k [(\mathbf{P}_k - \mathbb{E} \mathbf{P}_k)^2 + \mathbb{E}(\mathbf{P}_k - \mathbb{E} \mathbf{P}_k)^2] \\ &\preceq \frac{1}{2} \sum_k [2\mathbf{P}_k^2 + 2(\mathbb{E} \mathbf{P}_k)^2 + \mathbb{E} \mathbf{P}_k^2 - (\mathbb{E} \mathbf{P}_k)^2] \\ &\preceq \sum_k (\mathbf{P}_k^2 + \mathbb{E} \mathbf{P}_k^2). \end{aligned}$$

The first inequality follows from the operator convexity (1.1) of the square function; we have computed the second expectation exactly. The last bound follows from the operator Jensen inequality (1.2). Now, the matrix BDG inequality yields

$$\begin{aligned} E^2 &\leq \sqrt{2p-1} \cdot \left(\mathbb{E} \|\Delta_{\mathbf{X}}\|_p^p \right)^{1/2p} + \mu \\ &\leq \sqrt{2p-1} \cdot \left(\mathbb{E} \left\| \sum_k (\mathbf{P}_k^2 + \mathbb{E} \mathbf{P}_k^2) \right\|_p^p \right)^{1/2p} + \mu \\ &\leq \sqrt{4p-2} \cdot \left(\mathbb{E} \left\| \sum_k \mathbf{P}_k^2 \right\|_p^p \right)^{1/2p} + \mu. \end{aligned}$$

The third line follows from the triangle inequality for the L_p norm and Jensen's inequality.

Next, we search for a copy of E^2 inside this expectation. To accomplish this goal, we want to draw a factor \mathbf{P}_k off of each term in the sum. The following result of Pisier and Xu [PX97, Lem. 2.6] has the form we desire.

Proposition B.1 (A Matrix Schwarz-Type Inequality). *Consider a finite sequence $(\mathbf{A}_k)_{k \geq 1}$ of deterministic psd matrices. For each $p \geq 1$,*

$$\left\| \sum_k \mathbf{A}_k^2 \right\|_p \leq \left(\sum_k \|\mathbf{A}_k\|_{2p}^{2p} \right)^{1/2p} \left\| \sum_k \mathbf{A}_k \right\|_{2p}$$

Apply the matrix Schwarz-type inequality, Proposition B.1, to reach

$$\begin{aligned} E^2 &\leq \sqrt{4p-2} \cdot \left[\mathbb{E} \left(\sum_k \|\mathbf{P}_k\|_{2p}^{2p} \right)^{1/2} \left\| \sum_k \mathbf{P}_k \right\|_{2p}^p \right]^{1/2p} + \mu \\ &\leq \sqrt{4p-2} \cdot \left(\sum_k \mathbb{E} \|\mathbf{P}_k\|_{2p}^{2p} \right)^{1/4p} \left(\mathbb{E} \left\| \sum_k \mathbf{P}_k \right\|_{2p}^{2p} \right)^{1/4p} + \mu. \end{aligned}$$

The second bound is the Cauchy–Schwarz inequality for expectation. The resulting estimate takes the form $E^2 \leq cE + \mu$. Solutions of this quadratic inequality must satisfy $E \leq c + \sqrt{\mu}$. We reach

$$E \leq \sqrt{4p-2} \cdot \left(\sum_k \mathbb{E} \|\mathbf{P}_k\|_{2p}^{2p} \right)^{1/4p} + \left\| \sum_k \mathbb{E} \mathbf{P}_k \right\|_{2p}^{1/2}.$$

Square this expression to complete the proof of (7.3).

B.2. A Sum of Centered, Random Hermitian Matrices. We are now prepared to establish the bound (7.4) for a sum of centered, independent, random Hermitian matrices. Define the random matrix $\mathbf{X} := \sum_k \mathbf{Y}_k$. We may use the matrix Stein pair described in Section 2.4. According to (2.6), the conditional variance $\Delta_{\mathbf{X}}$ takes the form

$$\Delta_{\mathbf{X}} = \frac{1}{2} \sum_k (\mathbf{Y}_k^2 + \mathbb{E} \mathbf{Y}_k^2).$$

The matrix BDG inequality, Theorem 7.1, yields

$$\begin{aligned} \left(\mathbb{E} \|\mathbf{X}\|_{4p}^{4p} \right)^{1/4p} &\leq \sqrt{4p-1} \cdot \left(\mathbb{E} \|\Delta_{\mathbf{X}}\|_{2p}^{2p} \right)^{1/4p} \\ &= \sqrt{4p-1} \cdot \left(\mathbb{E} \left\| \frac{1}{2} \sum_k (\mathbf{Y}_k^2 + \mathbb{E} \mathbf{Y}_k^2) \right\|_{2p}^{2p} \right)^{1/4p} \\ &\leq \sqrt{4p-1} \cdot \left(\mathbb{E} \left\| \sum_k \mathbf{Y}_k^2 \right\|_{2p}^{2p} \right)^{1/4p}. \end{aligned}$$

The third line follows from the triangle inequality for the L_{2p} norm and Jensen’s inequality. To bound the remaining expectation, we simply note that the sum consists of independent, random psd matrices. We complete the proof by invoking the matrix Rosenthal inequality (7.3) and simplifying.

APPENDIX C. PROOF OF THEOREM 10.3: A COMBINATORIAL SUM OF MATRICES

Consider the matrix Stein pair $(\mathbf{X}, \mathbf{X}')$ constructed in Section 10.2. The expression (10.2) and the operator convexity (1.1) of the matrix square allow us to bound the conditional variance as follows.

$$\begin{aligned} \Delta_{\mathbf{X}}(\pi) &= \frac{1}{4n} \sum_{j,k=1}^n [\mathbf{A}_{j\pi(j)} + \mathbf{A}_{k\pi(k)} - \mathbf{A}_{j\pi(k)} - \mathbf{A}_{k\pi(j)}]^2 \\ &\preceq \frac{1}{n} \sum_{j,k=1}^n [\mathbf{A}_{j\pi(j)}^2 + \mathbf{A}_{k\pi(k)}^2 + \mathbf{A}_{j\pi(k)}^2 + \mathbf{A}_{k\pi(j)}^2] \\ &= 2 \sum_{j=1}^n \mathbf{A}_{j\pi(j)}^2 + \frac{2}{n} \sum_{j,k=1}^n \mathbf{A}_{jk}^2 \\ &= \mathbf{W} + 4\mathbf{\Sigma} \end{aligned}$$

where

$$\mathbf{W} := 2 \left(\sum_{j=1}^n \mathbf{A}_{j\pi(j)}^2 \right) - 2 \mathbf{\Sigma} \quad \text{and} \quad \mathbf{\Sigma} := \frac{1}{n} \sum_{j,k=1}^n \mathbf{A}_{jk}^2.$$

Substitute the bound for $\Delta_{\mathbf{X}}(\pi)$ into the definition (5.1) of $r(\psi)$ to see that

$$r(\psi) := \frac{1}{\psi} \log \mathbb{E} \bar{\text{tr}} e^{\psi \Delta_{\mathbf{X}}(\pi)} \leq \frac{1}{\psi} \log \mathbb{E} \bar{\text{tr}} e^{\psi(\mathbf{W} + 4\mathbf{\Sigma})} \leq 4\sigma^2 + \frac{1}{\psi} \log \mathbb{E} \bar{\text{tr}} e^{\psi \mathbf{W}}. \quad (\text{C.1})$$

The inequalities follow from the monotonicity of the trace exponential [Pet94, Sec. 2] and the fact that $\sigma^2 = \|\mathbf{\Sigma}\|$. Therefore, it suffices to bound the trace mgf of \mathbf{W} .

Our approach is to construct a matrix Stein pair for \mathbf{W} and to argue that the associated conditional variance $\Delta_{\mathbf{W}}(\pi)$ satisfies a semidefinite bound. We may then exploit the trace mgf bounds from Lemma 4.3. Observe that \mathbf{W} and \mathbf{X} take the same form: both have mean zero and share the structure of a combinatorial sum. Therefore, we can study the behavior of \mathbf{W} using the matrix Stein pair from Section 10.2. Adapting (10.2), we see that the conditional variance of \mathbf{W} satisfies

$$\begin{aligned} \Delta_{\mathbf{W}}(\pi) &= \frac{1}{n} \sum_{j,k=1}^n [\mathbf{A}_{j\pi(j)}^2 + \mathbf{A}_{k\pi(k)}^2 - \mathbf{A}_{j\pi(k)}^2 - \mathbf{A}_{k\pi(j)}^2]^2 \\ &\preceq \frac{4}{n} \sum_{j,k=1}^n [\mathbf{A}_{j\pi(j)}^4 + \mathbf{A}_{k\pi(k)}^4 + \mathbf{A}_{j\pi(k)}^4 + \mathbf{A}_{k\pi(j)}^4] \\ &\preceq \frac{4R^2}{n} \sum_{j,k=1}^n [\mathbf{A}_{j\pi(j)}^2 + \mathbf{A}_{k\pi(k)}^2 + \mathbf{A}_{j\pi(k)}^2 + \mathbf{A}_{k\pi(j)}^2]. \end{aligned}$$

In the first line, the centering terms in \mathbf{W} cancel each other out. Then we apply the operator convexity (1.1) of the matrix square and the bound $\mathbf{A}_{jk}^4 \preceq R^2 \mathbf{A}_{jk}^2$. Finally, identify \mathbf{W} and $\mathbf{\Sigma}$ to reach

$$\Delta_{\mathbf{W}}(\pi) \preceq 4R^2(\mathbf{W} + 4\mathbf{\Sigma}) \preceq 4R^2 \cdot \mathbf{W} + 16R^2\sigma^2 \cdot \mathbf{I}. \quad (\text{C.2})$$

The matrix inequality (C.2) gives us access to established trace mgf bounds. Indeed,

$$\log \mathbb{E} \bar{\text{tr}} e^{\psi \mathbf{W}} \leq \frac{8R^2\sigma^2\psi^2}{1 - 4R^2\psi}$$

as a consequence of Lemma 4.3 with parameters $c = 4R^2$ and $v = 16R^2\sigma^2$.

At last, we substitute the latter bound into (C.1) to discover that

$$r(\psi) \leq 4\sigma^2 + \frac{8R^2\sigma^2\psi}{1 - 4R^2\psi}.$$

In particular, setting $\psi = (8R^2)^{-1}$, we find that $r(\psi) \leq 6\sigma^2$. Apply Theorem 5.1 to wrap up.